arXiv:1512.01416v2 [cs.LO] 4 Nov 2016

New Lace and Arsenic: adventures in weak memory with a program logic (v2)

Richard Bornat, Middlesex University (R.Bornat@mdx.ac.uk)
Jade Alglave, Microsoft Research
Matthew Parkinson, Microsoft Research

Version 2, November 2016

Abstract

We describe a program logic for weak memory (also known as relaxed memory). The logic is based on Hoare logic within a thread, and rely/guarantee between threads. It is presented via examples, giving proofs of many weak-memory litmus tests. It extends to coherence but not yet to synchronised assignment (compare-and-swap, load-logical/store-conditional). It deals with conditionals and loops but not yet arrays or heap.

The logic uses a version of Hoare logic within threads, and a version of rely/guarantee between threads, with five stability rules to handle various kinds of parallelism (external, internal, propagation-free and two kinds of in-flight parallelism). There are $\mathbb{B}$ and $\mathbb{U}$ modalities to handle propagation, and temporal modalities since, $\mathbb{S}$ofar and $\mathbb{O}$uat to deal with global coherence (SC per location).

The logic is presented by example. Proofs and unproofs of about thirty weak-memory examples, including many litmus tests in various guises, are dealt with in detail. There is a proof of a version of the token ring.

A note on authorship This paper reports the joint work of its authors. But the words in the paper were written by Richard Bornat. Any opprobrium, bug reports, complaints, and observations about sins of commission or omission should be directed at him.

A note on version 2 The correspondence with Herding Cats has been clarified. The stability rules have been simplified: in particular the sat and $x{=}\widehat{x}$ tests have been eliminated from external stability checks. The embedding is simplified and has a more transparent relation to the mechanisms of the logic. Definitions of $\mathbb{U}$, $\mathbb{S}$ofar and $\mathbb{O}$uat have been considerably altered. The description of modalities and the treatment of termination has been reworked. Many proofs are reconstructed. A comprehensive summary of the logic is an appendix.

Contents

List of Definitions iv

List of Rules iv

I Background 1

1 Introduction 2

2 Choose Your Own Adventure 3

3 (S) (C) Notation 3

4 (S) Related work 6

4.1 Related work in program logics 6

5 (C) A logic for all seasonings 7

II Lacing, embroidery, modality 8

6 Sequentially-consistent concurrency 9

6.1 Interference and stability in SC 9

6.2 A checklist for proofs 10

6.3 Checking the proof of SC message-passing 12

7 The so tree and program order po 13

8 Weak-memory executions 13

9 Lacing to constrain weak memory, embroidery to make a proof 15

9.1 Embroidered constraints and the $\mathbb{B}$ modality 17

9.2 The same game on a new pitch 19

10 Checking the proof of weak-memory message passing 19

10.1 (C) Compiling laced programs 20

11 New instabilities 20

11.1 Internal (lo) parallelism 20

11.2 In-flight (bo) parallelism 23

11.3 Bo parallelism in the rely 25

11.4 Still the same game 25

12 Reads and the B modality 26

12.1 Bo from read to write 27

12.2 The B modality transmitted by interference 28

13 Auxiliary information in composite writes **29**

14 The lo, bo fragment **32**

15 Structured commands and disjunction of knots **32**
15.1 Conditionals 32
15.2 The constraint-coverage principle 33
15.3 Proof of the conditional receiver 33
15.4 Conditional lacing summary 34
15.5 (C) Compiling constraints with a control-expression source 34
15.6 Loops 35

III More precision and more orderings **38**

16 Interference and stability **39**
16.1 Internal and external parallelism 39
16.2 In-flight interference 40
16.3 A summary of stability checks 41

17 Implicit constraints **41**
17.1 SCloc (global coherence) 42
17.1.1 (P) RDW and CoRR litmus tests 42
17.1.2 A retreat 43
17.2 Register renaming (SCreg) 43
17.3 Quotienting internal interference 44
17.4 An example that appeals to SCreg dependencies 44

18 Speculated propagation and the go ordering **45**
18.1 (P) PPOCA: a scary conditional 47

IV Temporal reasoning and coherence **49**

19 A proof of SB using $\mathbb{U}$ and since **50**

20 The $\mathbb{S}$ofar modality **53**

21 Interference without propagation – UEXT instability **53**

22 In-flight UO instability **55**

23 A proof of LB using $\mathbb{S}$ofar **56**

24 The $\mathbb{O}$uat modality **58**

25 Reasoning about coherence order **59**
25.1 A first coherence example 59
25.2 Coherence and $\mathbb{B}$ 61
25.3 2+2W 62

26 SCloc is not optional **65**
26.1 Coherence from unrepeatable writes 66

V Termination and a loose end **67**

27 $\downarrow P$, the propagatable version of P **68**

28 A proper treatment of $\mathbb{B}$ and $\mathbb{U}$ **69**

29 A proper treatment of termination **70**
29.1 R: a problematic litmus test 70
29.2 The termination problem solved 73

VI A worked example **74**

30 The token ring **75**
30.1 A single-latch, single-token ring 78
30.2 A multi-latch, multi-token-variable ring 79
30.3 Impossible so paths 83

VII And finally **84**

31 Auxiliary variables and registers **85**

32 Inclusion of interference in guarantee (or rely) **86**

33 Library Abstraction and Optimisation **86**

34 Deficiencies **87**

References **88**

A Summary of the logic **93**

B (S) An SMT embedding **98**

List of Definitions

1 Postcondition of a component . . . 11
2 Strongest postcondition sp . . . 11
3 Substitution and the $\mathbb{B}$ modality . . . 18
4 Properties of the $\mathbb{B}$ modality . . . 18
5 Lo parallelism . . . 22
6 Bo parallelism . . . 24
7 Constraint-coverage principle . . . 33
8 Postcondition of a control expression . . . 34
9 Iterated constraint . . . 36
10 Hatting . . . 40
11 Precondition from knot . . . 46
12 Additional properties of the $\mathbb{U}$ modality . . . 51
13 Substitution and hatting of the $\mathbb{U}$ modality . . . 52
14 Properties of since . . . 52
15 Substitution and hatting of since . . . 52
16 Additional properties of the $\mathbb{S}$ofar modality . . . 53
17 Substitution and hatting of the $\mathbb{S}$ofar modality . . . 53
18 $\mathbb{B}$ and $(\tilde{})$ . . . 55
19 Uo parallelism . . . 55
20 Substitution and hatting of $\mathbb{O}$uat . . . 58
21 Properties of $\mathbb{O}$uat . . . 58
22 Coherence axioms . . . 60
23 Coherence hooking and hatting . . . 61
24 $\downarrow P$ – positive occurrences – and $\Downarrow P$ – negative occurrences . . . 69
25 Restricted $\mathbb{B}$, $\mathbb{U}$ and initial assertion . . . 69
26 The PMS assertion . . . 73
27 PMS coherence axiom . . . 73
28 Effect of a single interference . . . 86
29 Strongest post sp in the presence of temporal assertions . . . 99

List of Rules

1 SC stability . . . 11
2 Inheritance of embroidery . . . 17
3 External stability against $\mathbb{B}$-modal interference . . . 18
4 LO stability . . . 22
5 BO stability with $\mathbb{B}$-modal interference . . . 24
6 Stability against quotiented interference . . . 27

7 EXT stability 39
8 In-flight (bo) stability 40
9 UEXT stability 54
10 UO stability 55
11 Inclusion of interference 86

Part I
Background

1 Introduction

But you see ... **Well, insanity runs in my family. It practically gallops.** Concurrent programming is in a weird place. We write our programs in sequential languages (e.g. C (C, 2011), C++ (C++, 2014) or Java (Gosling et al., 2015)) but modern hardware (e.g. ARM (ARM, 2010), IBM Power (Pow, 2009), Intel x86 (x86, 2009) and any GPU) is mostly parallel: commands may not be carried out in the order they are written, and writes to memory may be substantially reordered. The consequences of parallelism are hidden from the naive user: ARM, Power and x86, for example, each ensure that a single isolated thread[1] – one which doesn't share memory with other threads – is the same initial state to final state mechanism that it would be on sequential hardware. But threads which share memory can observe the non-sequential behaviour of each other.

The effects of instruction-execution reordering and reordering of writes to memory are called *weak memory* or *relaxed execution*. We say weak memory.

Oh darling, just because Teddy's a little strange, that doesn't mean ... To write a shared-memory concurent program that is to run on weak-memory hardware, a programmer must understand where reordering needs to be constrained, and know how to constrain it with special instructions called *fences* or *barriers*, and/or with special assembly-code instruction-ordering tricks. There's a persuasive argument, put forward in (Boehm and Adve, 2012), that shared-memory concurrency is so complicated that users should remain ignorant and leave the difficult work to the experts, but those who write concurrency libraries, define languages or write compilers need formal tool support to ensure that they are getting the constraints right.

In the absence of formal models for concurrent systems, concurrent programs are often described as if they were to execute on a sequentially-consistent machine (Lamport, 1979), i.e. as if their executions were interleavings of the executions of separate threads. To execute those programs on weak-memory hardware we need to know whether and how to constrain the hardware, and how to do so on different architectures and on different implementations of those architectures.

Our focus is on programming, and on verification of programs. We worked for three years on a program logic for weak memory, which we called Lace logic. We targeted IBM Power, because it is one of the weakest and one of the most studied of modern architectures. But our logic ended up too complicated to be palatable, and with too many corner-cases for us to believe that a proof of soundness was possible.

Nevertheless we learnt a great deal about how to reason about programs under weak memory. Inspired by Crary and Sullivan (2015) we realised belatedly that we might make a simpler logic which deals with a wide variety of weak-memory architectures. This paper presents our first steps in that direction; this version retraces and revises those initial steps.

To back up, and somewhat validate, the work presented in this paper, we have a proof checker called Arsenic (Ars, 2015) and tens of worked examples. All our example proofs have been mechanically transliterated from mechanically-checked proofs (sometimes LaTeX-tweaked a little to fit on the page).

Another milestone would be a proof that our New Lace logic is sound with respect to existing models of concurrent systems, such as those put forward by (Alglave et al., 2014). We are not there yet; the concrete boots of soundness remain unfilled. One of our difficulties is that our logic is *thread-local* or *compositional* – a program proof is made up of separate proofs of its threads – whereas those modals deal with *global*

[1] As is common in the weak-memory research community, we use 'thread' to mean 'hardware thread', the smallest unit of processing supported by hardware. A single processor may have many cores, each of which may support several threads, each thread with its own registers.

properties of program executions. A soundness proof of our logic would likely have two stages: a proof of soundness against a compositional model of thread execution, and a proof of soundness of that model against global models of hardware. We have not yet attempted either part of that exercise.

2 Choose Your Own Adventure

Weak memory has surprising properties, and to deal with it our logic must be innovative. We expect that most of our readership will have a reasonable familiarity with Hoare logic, but not so many will have experience of concurrency proofs and rely/guarantee, and very few will understand weak memory. We have decided to present through examples, explaining as we go, rather than to present logic first and examples later. Some topics are dealt with incrementally, with forward references to the complete description later in the paper; it's not necessary to follow those references on a first reading.

We cover a great deal of ground over very many pages, and we expect that our readers may wish, at first, to skip some of our exposition. Our main purpose is to say what our logic is, via worked examples. We also have to justify our design, to explain the semantic background, show how it relates to existing hardware, and explain how we deal with some of the difficult problems which are well-known to researchers in the area. To help you navigate, we have labelled our sections and subsections with particular marks.

- Unmarked sections are mainstream;
- (C): how to relate laced programs to machine-code;
- (P): how we deal with well-known problematic examples;
- (S): questions of semantics.

For the brave and knowledgeable, there is a summary of our definitions and rules in appendix A. There is also appendix B which explains the embedding of our logic into the Z3 SMT engine to produce the Arsenic proof-checker.

3 (S) (C) Notation

Our programs are written in a notation given in table 1. Our notation is inspired by that of separation logic (Reynolds, 2002), with local registers and shared variables, and assuming no aliasing.[2] Programs manipulate integers and Booleans. Access to memory is only in assignment commands, and each assignment makes at most one memory access. There are various forms of auxiliary assignments.

Each *expr*, in an assignment or a control expression, is entirely local, referring only to non-auxiliary registers, constants and logical variables. Auxiliary *auxexpr*s may include auxiliary variables, but may only be assigned to auxiliary registers and variables. Tuples consisting of an *expr* and one or more *auxexpr*s may be assigned to a regular variable; the value assigned may be read simultaneously into a regular register and one or more auxiliary registers. A regular variable and one or more auxiliary variables may be assigned simultaneously.

Primitive commands – skips, assignments, asserts and control expressions – are labelled; labels are unique within a thread. There are no declarations. Programs are usually preceded by a labelled initial assertion and followed by a labelled final assertion.

Ours is not a high-level language, and in this paper we do not consider optimisations that might be attempted by a compiler. We note that optimisation in weak memory is problematic (see section 33). What we have is

[2] Though inspired by separation logic to concentrate on single writes, we do not yet approach the problem of the heap.

Table 1: programming language

program	::=	*passert* (*threads*) *passert*
passert	::=	'{' *label* : *assertion* '}'
threads	::=	*thread*
	|	*thread* '||' *threads*
thread	::=	*seq*
seq	::=	*command*
	|	*command* ; *seq*
	|	
command	::=	*label* : skip
	|	*label* : *assign*
	|	*label* : assert *assertion*
	|	if *label* : *expr* then *seq* fi
	|	if *label* : *expr* then *seq* else *seq* fi
	|	while *label* : *expr* do *seq* od
	|	do *seq* until *label* : *expr*
assign	::=	*write*
	|	*read*
	|	*calculation*
write	::=	*var* := *expr*
	|	*var* := *expr* [, *auxexpr*]$^+$
	|	*var* [, *auxvar*]^n := *expr* [, *auxexpr*]^n
	|	*auxvar* := *auxexpr*
read	::=	*reg* := *var*
	|	*reg* [, *auxreg*]$^+$:= *var*
calculation	::=	*reg* := *expr*
	|	*auxreg* := *auxexpr*
reg	::=	any name starting with 'r', but not starting with "raux"
auxreg	::=	any name starting with "raux"
label	::=	any name
var	::=	any name starting with a lower-case letter other than 'r', but not starting with "aux"
auxvar	::=	any name starting with "aux"
logicalvar	::=	any name starting with an upper-case letter
expr	::=	any expression whose operands are regs, constants (integers, true, false) and logvars, and whose operators are arithmetic or Boolean-arithmetic ($\neg, \wedge, \vee, \Rightarrow, \Leftrightarrow$)
expr	::=	as expr, but allowing auxregs
assertion	::=	any predicate-calculus expression; may include modalities $\mathbb{B}$, $\mathbb{U}$, $\mathbb{S}$ofar, $\mathbb{O}$uat, since and coherence assertions $var_c(expr, expr)$

Table 2: Laced threads

thread	::=	*guarantee seq threadpost rely*
guarantee	::=	guar '[' *interferences* ']'
rely	::=	rely '[' *interferences* ']'
	|	
interferences	::=	*interference*
	|	*interference* ; *interferences*
	|	
interference	::=	*assertion* '|' *var* := *expr*
	|	'[' *names* ']' . *assertion* '|' *var* := *expr*
names	::=	*name*
	|	*name* , *names*
threadpost	::=	*simpleknot*
	|	
simpleknot	::=	'{' *stitches* '}'
	|	*simpleknot* $\vee$ *simpleknot*
	|	*simpleknot* $\overrightarrow{\vee}$ *simpleknot*
knot	::=	*simpleknot intfpre*
	|	
intfpre	::=	'[' *assertion* ']'
	|	
stitches	::=	*stitch*
	|	*stitch* ; *stitches*
	|	
stitch	::=	*labref ordering sourcepost* : *assertion*
labref	::=	*label*
	|	$label_t$
	|	$label_f$
ordering	::=	lo
	|	bo
	|	uo
	|	go
sourcepost	::=	'{' *assertion* '}'
	|	
command	::=	*knot label* : skip
	|	*knot label* : *assign*
	|	*knot label* : assert *assertion*
	|	if *knot label* : *expr* then *seq* fi
	|	if *knot label* : *expr* then *seq* else *seq* fi
	|	while *knot label* : *expr* do *seq* od
	|	do *seq* until *knot label* : *expr*

close to the machine, a sort of architecture-independent assembly code with structured commands.

In introducing our examples we deal with the annotations that are necessary for programming, introducing them incrementally and explaining their meanings. For reference we give in table 2 a definition of laced threads – threads adorned with constraints for weak memory, with an interference guarantee and an optional interference rely. It is not necessary on first reading to study this table, nor at first to understand the concepts of constraint, ordering, interference, guarantee and rely. Some of the non-terminals in table 2 take their definitions from table 1.

4 (S) Related work

A very great deal of work has been expended on models of weak memory, compilation for weak memory and verification of programs in languages designed to work on weak memory. We do not attempt a complete bibliography, but we list work which is immediately relevant to ours.

- On models of hardware: Alglave et al. (2009), Sarkar et al. (2009), Owens et al. (2009), Sewell et al. (2010). Alglave et al. (2011), Sarkar et al. (2009, 2011), Alglave and Maranget (2011), Alglave et al. (2012), Mador-Haim et al. (2012), Alglave et al. (2014, 2015).
- On the semantics of programming languages for weak memory: Boehm and Adve (2008), Batty et al. (2011), Boehm and Adve (2012), Boehm and Demsky (2014), Sarkar et al. (2012), Vafeiadis et al. (2015), (Vafeiadis, 2015).
- On verification of programs in weak-memory-savvy languages: Batty et al. (2013), Lahav and Vafeiadis (2015), Tassarotti et al. (2015), Turon et al. (2014), Vafeiadis and Narayan (2013), Norris and Demsky (2013), Alglave et al. (2013), Burckhardt and Musuvathi (2008), Atig et al. (2010), Atig et al. (2012).
- On compilation for weak memory: Kuperstein et al. (2010), Linden and P.Wolper (2011), Sevcík et al. (2011), Vafeiadis and Zappa Nardelli (2011), Atig et al. (2011), Abdulla et al. (2012), Abdulla et al. (2013).

Our work relates particularly to the work on models of hardware. We take logical reasoning techniques for SC concurrency developed by Owicki (1975), Owicki and Gries (1976a,b) and Jones (1983), and, using intuitions derived from those models and the experiments which underpin them, port those techniques to weak memory.

We hope that we'll be able in the future to add various features to our logic, as was done for SC concurrency with Separation Logic (Reynolds, 2002; Calcagno et al., 2006), Concurrent Abstract Predicates (Dinsdale-Young et al., 2010), and RGSep (Vafeiadis and Parkinson, 2007; Vafeiadis, 2007).

4.1 Related work in program logics

We note several attempts to reason about programs running on TSO hardware (x86, notably, is TSO) and to reason about C11 programs.

In reasoning about TSO, Ridge (2010) made a program logic in which the write buffers described by Owens et al. (2009) were dealt with explicitly. Wehrman and Berdine (2011) abstracted the frame buffers in a sub-structural logic which used a frame rule to handle the effect of delayed writes to memory. Sieczkowskie (Sieczkowski, 2013; Sieczkowski et al., 2015) abstracted the write buffers entirely in a temporal logic

Table 3: The orderings of lace logic and their equivalents in the cats model and in hardware

	cats	x86	Power	ARM
lo	ii/ci	implicit	implicit (register assign→use); implicit (same location, read→assign / assign→read / assign→assign); otherwise data (calc→calc), addr (read→read, read→assign) or rfi-addr (assign→assign, assign→read) or ctrl-isync (control-expr→read)	implicit (register assign→use); implicit (same location, assign→read); otherwise otherwise data (calc→calc), addr (read→read, read→assign) or rfi-addr (assign→assign, assign→read) or ctrl-isb (read→read)
bo	lwfence	implicit	lwsync	dsb
uo	ffence	MFENCE	sync	dsb
go	ic	implicit	ctrl	implicit (same location, read→assign or assign→assign); otherwise rfi-ctrl or ctrl

which employed a 'fiction of sequential consistency', building on the concurrent abstract predicates notion (Dinsdale-Young et al., 2010).

In reasoning about C11, Batty et al. (2013) showed that the possibility of 'satisfaction cycles' made abstraction of libraries which use relaxed atomics impossible (see also our discussion in section 18). Vafeiadis et al. (Vafeiadis and Zappa Nardelli, 2011; Vafeiadis and Narayan, 2013; Turon et al., 2014; Vafeiadis et al., 2015; Vafeiadis, 2015; Lahav and Vafeiadis, 2015; Tassarotti et al., 2015) have reasonably therefore focussed on the release-acquire fragment of C11. Their logic effectively covers the same area as the bo, lo fragment which we deal with in sections 9 to 16 of part II of this paper. But their reach is far beyond ours: they employ the 'ownership transfer' notion of separation to deal with release-acquire and are able then to make proofs of important algorithms, especially including RCU (McKenney, 2004). We envy their work.

5 (C) A logic for all seasonings

Our logic is based on programs explicitly laced with constraints, each imposing an ordering between commands or between control expressions and commands. Those orderings are introduced and explained in part II. Table 3 shows our understanding (not yet backed up by a soundness proof) of the correspondence of our orderings to the treatment in (Alglave et al., 2014) and to programs running on the three major processor architectures. We expect that there will be correspondences to GPU architectures as well, but we don't know them yet.

In various places we have referred to compilation of laced programs into machine code. We expect that the correspondences in table 3 would be the basis of any such compiler, but we don't expect it to be easy to make a compiler. In particular fence/barrier placement is a hard problem (Burckhardt et al., 2007; Huynh and Roychoudhury, 2007; Atig et al., 2011), but dealing with multiple constraints in precondition knots is also by no means straightforward.

Part II

Lacing, embroidery, modality

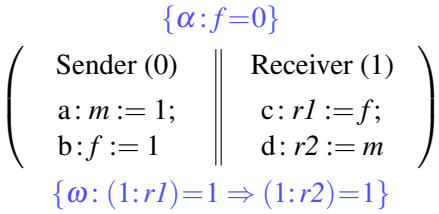

Figure 1: MP: shared-variable message-passing

6 Sequentially-consistent concurrency

Before we consider our treatment of weak memory, we explain the basics of our reasoning, using an example program which will be the basis of many later examples.

Sequentially-consistent – henceforth *SC* – execution (Lamport, 1979), is what machines did before weak memory upset the applecart. Within a thread, SC execution follows the *sequential order* – henceforth so – of commands. Within in a program, thread executions are *interleaved*. Because there is a single shared memory, each read $r := x$ in the interleaving takes its value from the latest preceding write $x := E$, no matter which thread made that write, or from the initial state, if there is no preceding write. If we prefix the interleaving with writes which establish the initial state, we can say uniformly that each read in the interleaving takes its value from the latest preceding write to the same variable.

Consider SC execution of the two-thread program in fig. 1. The initial-state assertion α[3] requires that execution begins with 0 in variable f. The sender (thread 0) writes 1 into message variable m and then writes 1 into flag variable f; the receiver (thread 1) reads from f and then from m. If the execution interleaving is 'α,a,b,c,d' then command c must read 1 from the write of command b, and d must read 1 from the write of a. In all other interleavings c precedes b and so will read 0 from α, and command d may or may not read 1, depending on whether or not it follows a. So we can only be certain that d reads 1 when c has already read 1. The final-state assertion ω states this formally, using the notation $(i{:}reg)$ for 'register *reg* in thread *i*'.

Informal reasoning about execution of this example is easy, but in more complex cases it is very difficult to consider all possible interleavings. We'd like to prove formally that every execution of the program which starts in a state satisfying precondition α will, *if* it terminates, do so in a state satisfying postcondition ω. We shan't be able to prove that every execution starting from α *must* terminate: even in such a simple example, our logic doesn't run to it.

6.1 Interference and stability in SC

Our SC-concurrency proof of MP is shown in fig. 2. There are state assertions enclosed in braces before, between and after the commands of each thread; in each triple $\{P\}\,C\,\{Q\}$ precondition P, command C and postcondition Q are formally related by the logic's command-execution rules. In multi-thread programs the assertions are further constrained, because one thread may change the values of shared variables used by another.

Interleaving, from the point of view of a thread, means that other threads' commands execute in the gaps before, between and after its own command executions, just where we place the proof assertions. We say

[3] In our logic the program precondition must be labelled, in order that we can refer to in when placing constraints on execution. By convention we use the label α. The program postcondition doesn't really need a label, but for symmetry we require the label ω.

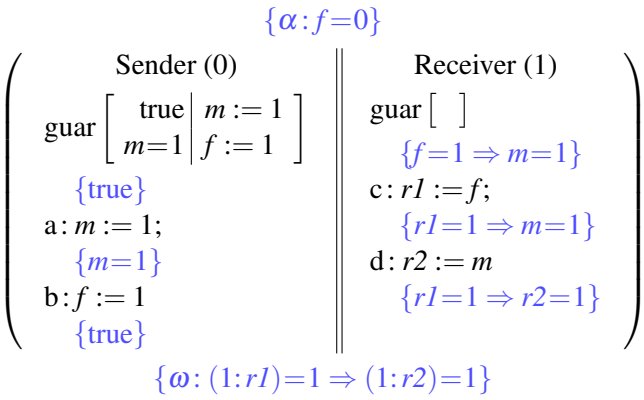

Figure 2: SC proof of MP

Table 4: Per-thread concurrency-proof checklist

1. Check that each assertion is implied by the postcondition of the immediately-preceding component.
2. Check that each assertion is stable under the interference of parallel assignments.
3. Check that the stable interference of each variable assignment is included in the thread's guarantee.

that those interpolated executions are *parallel* with those assertions, and *interfere* with them by assigning values to variables on which they may depend. We require that assertions are *stable* against all parallel command executions, which means that when the assertion is true, none of the interference of any possibly-interpolated command can make it false. Only the initial and final assertions (α and ω in fig. 2, for example) are automatically stable, because they hold when no thread is executing.

If we can invent stable assertions which also fit the command-execution rules then we have a proof of concurrent execution. That's tricky, as we shall see, but it's what we need for a shared-variable concurrency proof.

Those are the bare bones of the longest-established method for verifying concurrent programs, a modification of Hoare logic (Hoare, 1969) devised by Owicki and Gries (Owicki, 1975; Owicki and Gries, 1976a,b). Jones's rely/guarantee (Jones, 1983) is a later development, abstracting interference from the commands and assertions of a thread, designed as a refinement method for concurrent-program development. The word 'stability' is more modern still (Xu et al., 1997), supplanting the earlier 'interference-free'.

We take from Jones the ideas of a *guarantee* which summarises the interference made by a thread, and a *rely* which summarises the interference of other threads which it must withstand. Like Owicki and Gries we concentrate on verifying proofs and we focus on the interference of individual assignments.

6.2 A checklist for proofs

It's not enough simply to present a proof like that in fig. 2. We have to check that its assertions fit the command-execution rules and are stable against interference from other threads. Our checklist is in table 4: check 1 works on single commands within a thread; check 2 is about interference and stability; check 3 is about making a summary of a thread's interference, which simplifies stability reasoning in other threads.

In each thread of each of our proofs there is a summary of its interference in a *guarantee* clause, prefixed

'guar'. In an SC guarantee we write $Q \mid x := E$ to describe the interference of an assignment $x := E$ with precondition Q: that description tells us that when Q holds the guaranteeing thread may execute $x := E$. The guarantee of the sender in fig. 2 claims that it makes two kinds of interference; the empty guarantee of the receiver claims that it doesn't make any interference at all. Each of those claims must be checked.

For check 1 we need to calculate postconditions of components from their preconditions. In this example the only components are assignments and the initial assertion. (In later examples the control expressions of conditionals and loops are also components and have pre and postconditions: see section 15.)

Definition 1: Postcondition of a component

The postcondition of the initial-state assertion is the assertion itself. The postcondition of skip is its precondition. The postcondition of assert P is the conjunction of its precondition and P. For assignments we use *strongest postconditions* (Dijkstra, 1976).

Strongest postconditions are given by Floyd's assignment rule (Floyd, 1967). We adapt it to the various forms of our assignments: a write $x := E$; a calculation $r := E$; and a read $r := x$.

Definition 2: Strongest postcondition sp

$$\begin{array}{lcl} \mathrm{sp}(P,\ x := E) & \triangleq & P[x \backslash \overleftarrow{x}] \wedge x = E \\ \mathrm{sp}(P,\ r := E) & \triangleq & P[r \backslash \overleftarrow{r}] \wedge r = E[r \backslash \overleftarrow{r}] \\ \mathrm{sp}(P,\ r := x) & \triangleq & P[r \backslash \overleftarrow{r}] \wedge r = x \end{array}$$

– $\overleftarrow{x}$ is the pre-assignment value of x and $\overleftarrow{r}$ the pre-assignment value of r. Floyd used existential quantification to quotient references to pre-assignment values, but since $\overleftarrow{x}$ and $\overleftarrow{r}$ can't appear in our proof assertions we don't need to quotient. Note that according to the syntax of our language in table 1 there can be no occurrences of x in E or r in x, which means that we can avoid substituting in either in the first and third lines of the definition.

Interference in SC execution happens when an assertion already holds and a variable assignment from another thread muscles in, changing the values of variables which underpin it.[4] Registers are private to a thread, so register assignments don't interfere with other threads' assertions.

Rule 1: SC stability

In SC execution, P is stable against $Q \mid x := E$ if

$$\mathrm{sp}(P \wedge Q,\ x := E) \Rightarrow P$$

– suppose P holds; suppose Q holds as well, so $x := E$ may be executed; we require that if it is executed, then P still holds afterwards.

There's a circularity in the notion of stability: thread A's assertions must be stable against the interference of thread B, using B's precondition-assertions in the stability check, and B's assertions must be stable against the interference of A, using A's precondition-assertions. This is quite ok semantically, and we shan't worry about its subtleties. Circularity can simplify proofs, so we make use of it.

[4] In this paper we deal only with registers and shared variables. In the future we hope to deal also with the heap, building perhaps on the techniques of RGSep (Vafeiadis and Parkinson, 2007; Vafeiadis, 2007).

6.3 Checking the proof of SC message-passing

We begin our check in the sender of fig. 2. Before command a executes, true certainly holds, so we can use it as a precondition. Following our checklist we note: the initial state α implies true (check 1); there isn't any interference, because the receiver's guarantee is empty (check 2); there's an entry in the sender's guarantee with a's precondition true and a's command $m := 1$ (check 3).

To deal with the precondition of command b, we must calculate the postcondition of a. Definition 2 gives the strongest postcondition true $\wedge\, m{=}1$, which implies $m{=}1$, the precondition we've used for b (check 1); stability is immediate, as before (check 2); there's an entry in the guarantee with b's command and precondition (check 3). And that's it for the sender: we've checked all its assertions with checks 1 and 2; we've checked all its interference with check 3.

The sender writes once to m and once to f, but its guarantee doesn't say that. From the guarantee it appears that it might interfere repeatedly and at any time with $m := 1$, and it might interfere repeatedly with $f := 1$ whenever $m{=}1$. We shall see that the receiver's assertions are stable despite this overstatement of interference. Sometimes it's useful to give more constraining interference preconditions, and to make sure that interference occurs only once, but not in this example.

The clever bit of our proof is the choice of $f{=}1 \Rightarrow m{=}1$ as precondition for command c in the receiver. Informally, considering the execution of the sender, when f is 1, m must already be 1. Formally, we must show (check 1) that the precondition of c is implied by the initial state: that's easy because the initial state gives $f{=}0$, so the precondition is trivially satisfied. Then we have to show (check 2) that the assertion is stable against the guaranteee of the sender. Rule 1 generates two checks, one for each of the sender's interferences. The first check is passed because the interference produces $m{=}1$:

$$\begin{array}{ll} & \mathrm{sp}((f{=}1 \Rightarrow m{=}1) \wedge \mathrm{true},\ m := 1) \\ = & (f{=}1 \Rightarrow \overleftarrow{m}{=}1) \wedge \mathrm{true} \wedge m{=}1 \\ \Rightarrow & m{=}1 \\ \Rightarrow & f{=}1 \Rightarrow m{=}1 \end{array} \tag{1}$$

and the second because the interference can only happen when $m{=}1$ holds already:

$$\begin{array}{ll} & \mathrm{sp}((f{=}1 \Rightarrow m{=}1) \wedge m{=}1,\ f := 1) \\ = & (\overleftarrow{f}{=}1 \Rightarrow m{=}1) \wedge m{=}1 \wedge f{=}1 \\ \Rightarrow & m{=}1 \\ \Rightarrow & f{=}1 \Rightarrow m{=}1 \end{array} \tag{2}$$

We don't have to put anything in the receiver's guarantee because c assigns to a register (check 3).

The strongest postcondition of command c, by definition 2, is $(f{=}1 \Rightarrow m{=}1) \wedge r1{=}f$, which implies $r1{=}1 \Rightarrow m{=}1$, the precondition of d (check 1); that in turn is stable against the first line of the sender's guarantee because $m{=}1$ is stable, and against the second line because it doesn't mention f (check 2); there's nothing to put in the guarantee for a register assignment (check 3).

We need a postcondition for the receiver, to justify the final assertion ω. The postcondition of d is $(r1{=}1 \Rightarrow m{=}1) \wedge r2{=}m$, which implies $r1{=}1 \Rightarrow r2{=}1$, which is stable because it doesn't mention any variables at all. And that's it for the receiver: we've checked the assertions with checks 1 and 2; there weren't any variable assignments for check 3.

The final assertion ω is derived from the conjunction of the postconditions of the two threads, annotated where necessary to make it clear which thread's registers we're talking about. Informally, we use just the receiver's postcondition, because the sender's postcondition true contributes nothing.

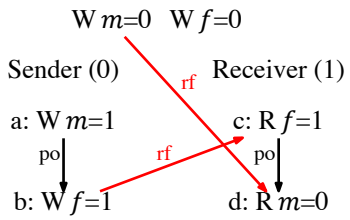

Figure 3: MP litmus execution

We have proved that under SC execution, and provided that it terminates, MP will do just what we think it ought to. We've been through the proof in great detail to show how the method works. Weak-memory execution introduces interesting new kinds of ordering and parallelism, but we shall use the same method, and the same checklist, to check our weak-memory proofs.

7 The so tree and program order po

In straight-line threads like those of fig. 1, sequential order so is just textual order. In more complicated cases conditionals must be expanded and loops unrolled to generate a potentially infinite so tree of command instances (see sections 15.1 and 15.6).

Every SC execution of a thread must follow a path in the so tree, but there may be many executions that follow the same path, differing according to the values in their registers and variables. *Program order* po orders command-instance *executions* within a thread execution – i.e. po is so with values.

8 Weak-memory executions

Weak-memory execution of a single thread is constrained to follow a *reordering* of a path in the thread's so tree. In our logic there are no other implicit constraints on the execution of a thread. That means that there can be significant parallelism between command executions even within a single thread, and the job of the programmer is to describe explicitly what constraints are needed to define the (usually partial) order they require.

Weak memory doesn't rely on interleaving of executions to give values to reads. Instead writes are *propagated* between threads – but not necessarily in any particular order and not necessarily to all threads in the same order or at the same time. There is no longer a global interleaving or a single global shared-memory state, and at any instant a thread may not yet have been affected by all the writes made by other threads. In our logic we suppose that propagation isn't entirely unconstrained, but the issue is a little subtle, and we defer a full discussion until section 17.1. Most of the time it's reasonable to suppose that there are no implicit constraints on the reordering of propagation.

So weak memory is not SC, and we illustrate that by presenting three classic *litmus-test* examples (Alglave et al., 2011). In each case a program which works under SC is observed to execute anomalously in weak memory. The first example is MP: on Power, on ARM and on Nvidia GPUs, experiment shows that the receiver can read 0 from m even though it reads 1 from f. The offending execution is pictured in fig. 3. The nodes are writes (W) and reads (R); the first row shows the writes that create the initial state (for simplicity, litmus tests initialise all variables); the two columns show the write and read actions of the threads; and the diagram explicitly shows rf arrows from write to read. Note that the diagram shows po ordering in each thread, but remember that the hardware takes little account of po.

$$\{\alpha : x{=}y{=}0\}$$
$$\left(\begin{array}{c|c} \text{Thread 0} & \text{Thread 1} \\ \text{a}: r1 := y; & \text{c}: r1 := x; \\ \text{b}: x := 1 & \text{d}: y := 1 \end{array}\right)$$
$$\{\omega : \neg((0{:}r1){=}1 \wedge (1{:}r1){=}1)\}$$

Figure 4: LB: a race with read before write

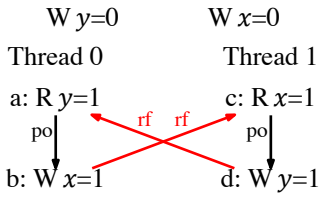

Figure 5: LB litmus execution

The diagram shows one *possible* weak-memory execution of MP, which happens sometimes but not always (and that makes it, in John Reynolds' words, the worst kind of bug). It doesn't happen under SC, as we proved it couldn't and as generations of programmers who have used the algorithm know very well. But in weak memory the command executions of the sender can be reordered so that f is written before m; the command executions of the receiver can be reordered to read m before f; the sender's f write can be propagated to the receiver before its m write; or there can be some combination of some or all or none of those things. To make the program work reliably in weak memory we have to impose constraints on execution and propagation reordering, described in the next section.

Our second example is LB, in fig. 4. Each thread reads from a variable – x in thread 0, y in thread 1 – so-before it writes to the other variable. In SC it would be impossible that they could each read from the other's write: command c can only read 1 in an 'α,a,b,c,d' interleaving, in which a must read 0; vice-versa a can only read 1 in an 'α,c,d,a,b' interleaving, in which c must read 0. But in weak memory it can and does happen, just because command executions can be reordered within the threads. Fig. 5 pictures the offending execution, which is observed on Power, ARM and Nvidia GPUs (at least). We deal with the problem in section 23.

Our third example is SB, in fig. 6. This time each thread writes to one variable – x in thread 0, y in thread 1 – so-before reading from the other variable. In any SC interleaving one thread's write has to come first, and

$$\{\alpha : x{=}0 \wedge y{=}0\}$$
$$\left(\begin{array}{c|c} \text{Thread 0} & \text{Thread 1} \\ \text{a}: x := 1; & \text{c}: y := 1; \\ \text{b}: r1 := y & \text{d}: r1 := x \end{array}\right)$$
$$\{\omega : \neg((0{:}r1){=}0 \wedge (1{:}r1){=}0)\}$$

Figure 6: SB: a race with write before read

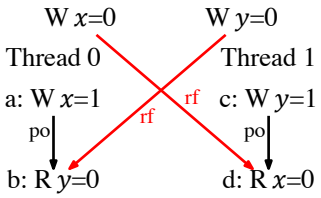

Figure 7: SB litmus excution

then the other thread's read must see 1. This code is part of Dekker's mutual-exclusion algorithm (Dijkstra, n.d.). In weak memory both threads can read 0, because of execution and/or propagation delays (even if each thread executes its commands in order, each may execute its read before the other's write has managed to reach it). The execution in fig. 7 is observed on several architectures, including x86, Power and ARM. We show how we deal with the problem in section 19.

9 Lacing to constrain weak memory, embroidery to make a proof

Weak memory, as we've described it, has very few constraints on reordering of execution and propagation. Any particular architecture may impose constraints of its own, but a programmer often needs to impose still more constraints by using special architecture-dependent instructions (barriers, fences) and/or coding devices (address and data dependencies). Programming is tricky because the effects of reordering are subtle, the architectural constraints are subtle, and the effects of special instructions and coding devices are subtle.

Following Crary and Sullivan (2015) we abstract from machine-dependent constraints and programming devices and *lace* our programs with explicit architecture-independent constraints. Lace constraints can then be translated into programming devices for a particular architecture, or ignored if the architecture already imposes that constraint.

Just like the architecture-dependent programming devices, lacing applies only within a thread, and only between components that are already so ordered. Lacing strengthens the partial order of execution and/or propagation of the thread.

Lace logic's constraints apply between labelled components of a thread, and they are ordered – a→b, c→d, etc. – each with a *source* and a *target*. In straight-line threads like the MP sender and receiver, labelled components are the initial assertion and primitive commands, but we shall see that we also label control expressions in conditionals and loops.

There are four kinds of constraint – lo, bo, uo and go – each imposing a particular kind of ordering. In laced programs each target component has a *knot* of constraints, enclosed in braces like a Hoare-logic precondition, each constraint giving an ordering and a source. In principle a knot may include several constraints, but in our first examples each will have only one.

Weak-memory execution is a tricky concept because commands don't 'execute' all at once. Instead we talk of *elaboration* of components and *propagation* of writes. Elaboration is whatever is required to establish a component's postcondition in a thread's local state; we suppose that elaboration completes all at once. Variable assignments create writes, and for a variable assignment propagation follows elaboration; we don't say that it follows immediately, and we don't suppose that propagation completes all at once in every destination thread.

An *elaboration order* lo constraint ensures that its source elaborates before its target; as a consequence the source's postcondition is available to the precondition of the target. Lo is the least we can ask for in order

$$\{\alpha : f{=}0\}$$
$$\left(\begin{array}{c|c} \text{Sender (0)} & \text{Receiver (1)} \\ \{\mathsf{lo}\,\alpha\}\ \mathrm{a} : m := 1; & \{\mathsf{lo}\,\alpha\}\ \mathrm{c} : r1 := f; \\ \{\mathsf{bo}\,\mathrm{a}\}\ \mathrm{b} : f := 1 & \{\mathsf{lo}\,\mathrm{c}\}\ \mathrm{d} : r2 := m \end{array}\right)$$
$$\{\omega : (1{:}r1){=}1 \Rightarrow (1{:}r2){=}1\}$$

Figure 8: MP laced for weak memory

$$\{\alpha : f{=}0\}$$
$$\left(\begin{array}{c|c} \text{Sender (0)} & \text{Receiver (1)} \\ \text{guar}\left[\begin{array}{r|l} \text{true} & m := 1 \\ \mathbb{B}(m{=}1) & f := 1 \end{array}\right] & \text{guar}\left[\ \right] \\ \{\mathsf{lo}\,\alpha : \text{true}\}\ \mathrm{a} : m := 1; & \{\mathsf{lo}\,\alpha : f{=}1 \Rightarrow m{=}1\}\ \mathrm{c} : r1 := f; \\ \{\mathsf{bo}\,\mathrm{a} : \mathbb{B}(m{=}1)\}\ \mathrm{b} : f := 1 & \{\mathsf{lo}\,\mathrm{c} : r1{=}1 \Rightarrow m{=}1\}\ \mathrm{d} : r2 := m \\ & \{\mathsf{lo}\,\mathrm{d} : r1{=}1 \Rightarrow r2{=}1\} \end{array}\right)$$
$$\{\omega : (1{:}r1){=}1 \Rightarrow (1{:}r2){=}1\}$$

Figure 9: A weak memory proof of MP

to be able to reason about states within a thread. In fig. 8, command a is lo-after the initial state α, and α, c and d are lo-ordered.

Lo doesn't control propagation: it can ensure that assignments are elaborated in so order, but their writes might still be propagated in another order, and they might be propagated in different orders to different destinations. A *before order* bo constraint controls both propagation and elaboration. Like lo it imposes elaboration order between source and target, and in addition it ensures that the writes which underpin the source's postcondition will be propagated to any thread before the writes made by the target.[5]

Bo is the least we can ask for in order to control propagation. In fig. 8, commands a and b are bo ordered, so they will be elaborated in so order, echoing SC execution, and the $m{=}1$ write will be propagated to the receiver before the $f{=}1$ write, echoing the effects of SC interleaving.

Universal order uo gives still more control over propagation, ensuring immediate propagation of the source-postcondition writes. It is discussed in section 19.

Propagation order go ensures that the source is elaborated before the target's writes are propagated. It is discussed in section 18.

Constraint orderings compose: $\langle \mathsf{lo};\mathsf{lo}\rangle$ is lo; both $\langle \mathsf{lo};\mathsf{bo}\rangle$ and $\langle \mathsf{bo};\mathsf{lo}\rangle$ are bo; each of $\langle \mathsf{lo};\mathsf{uo}\rangle$, $\langle \mathsf{bo};\mathsf{uo}\rangle$, $\langle \mathsf{uo};\mathsf{lo}\rangle$; $\langle \mathsf{uo};\mathsf{bo}\rangle$ is uo. Orderings nest: lo includes bo and uo; bo includes uo; go is on its own.

[5] This wording hides a rather complicated operational mechanism. Not all the writes will be propagated to all threads, but no source-postcondition write will be propagated to any particular thread after the target write is propagated to that same thread. In our discussion the distinction between 'in order' and 'not out of order' will not matter until section 29. Program logic, in this and in many other cases, abstracts and simplifies.

9.1 Embroidered constraints and the $\mathbb{B}$ modality

In Hoare-logic proofs for SC we write assertions before, between and after commands. In Lace logic weak-memory proofs we *embroider* constraints with assertions. Fig. 9 is an embroidered version of fig. 8; note the embroidered thread-postcondition knot after command d in the receiver. A novelty in this proof is the use of the $\mathbb{B}$ modality on the bo constraint a→b, and also in the guarantee entry generated by b. The $\mathbb{B}$ modality is explained below, but note that if we ignore the $\mathbb{B}$s then the assertions and the interference descriptions are just the same as those in the SC proof of fig. 2.

There is an important difference between SC assertions and weak-memory assertions. SC assertions are about the state of the global shared store, whereas weak-memory assertions are about a *thread-local view* of the store, more or less about the values in the asserting thread's cache. It's entirely possible, for example, that at some instant in a weak-memory execution of MP, even when constrained as in fig. 8, $m{=}1$ holds in the sender and $m{=}0$ in the receiver. That couldn't be true in SC.

The precondition of a command, in simple cases, is the conjunction of its knot's embroideries (for a fuller discussion, dealing with disjunction of knots and go constraints, see sections 15.1, 15.6 and 18). The postcondition of a component is derived, as in SC proofs, according to definition 1.

In an SC proof each assertion is implied by the preceding component's postcondition, and is in turn the precondition of the next component. In a weak-memory proof there's a lacier connection. The embroidery of a constraint depends on the postcondition of the constraint's source, but the relationship depends on the kind of constraint.

Rule 2: Inheritance of embroidery

If a constraint's source postcondition is P and the constraint's embroidery is Q then we require

- on an lo or go constraint, $P \Rightarrow Q$;
- on a bo constraint, for some R, $P \Rightarrow R$ and $\mathbb{B}(R) \Rightarrow Q$;
- on a uo constraint, for some R, $P \Rightarrow R$ and $\mathbb{U}(R) \Rightarrow Q$.

In many cases a bo constraint will be embroidered $\mathbb{B}(R)$ where $P \Rightarrow R$. A similar remark applies to uo constraints, but we defer further discussion of uo until section 19.

The embroidery on a constraint must be stable against external and (we shall see) internal interference. That's a weak-memory novelty: embroidery stability is per constraint, not per pre or postcondition. That's not the only kind of stability in weak memory: interference preconditions in the guarantee must also be checked for stability – but we'll come to that.

The $\mathbb{B}$ modality polices the propagation-ordering requirement of the bo constraint.

- $\mathbb{B}(P)$ in a precondition asserts that P has held continuously since $\mathbb{B}(P)$ was established in a bo constraint;
- $\mathbb{B}(Q) \mid x := E$ in a guarantee asserts that the writes which underpin Q will be propagated to an interfered-with thread before the interfering write $x{=}E$.

Aside 1: (S) (C) Whence the $\mathbb{B}$ modality?

Weak-memory hardware constrains propagation by special *barrier* instructions. A bo constraint between two commands will require, on sufficiently-weak hardware such as Power and ARM, a barrier instruction to be executed between those commands. The hardware will then ensure that the writes made or witnessed before the barrier will be propagated to any particular thread before those made after the barrier (more precisely, writes made after the barrier will not be propagated before those made or witnessed before).

Hoare logic puns between program variables and similarly-named logical variables. Lace logic puns also between assertions and writes. If the (stable) assertion P holds just before a barrier, then the thread has made or witnessed[6] the writes which gave variables the values implied by that precondition. So long as P continues to hold after the barrier we can use it to stand for those writes. $\mathbb{B}(P)$ on a constraint asserts that P has held *continuously* since a barrier was executed and therefore the writes which underpin it will be propagated before any write made by the target of the constraint. Note that $\mathbb{B}$ it is a particular instance of the since modality (section 19).

Because the P assertion must hold continuously since $\mathbb{B}(P)$ was established, $\mathbb{B}(P)$ has special treatment in substitution – i.e. in computing strongest postconditions according to definition 2.

Definition 3: Substitution and the $\mathbb{B}$ modality

$$\begin{array}{rcl} \mathbb{B}(P)[x\backslash\overleftarrow{x}] & = & \overleftarrow{\mathbb{B}(P)} \wedge P[x\backslash\overleftarrow{x}] \\ \overleftarrow{\mathbb{B}(P)} \wedge P & \Rightarrow & \mathbb{B}(P) \\ \mathbb{B}(P)[r\backslash\overleftarrow{r}] & = & \mathbb{B}(P[r\backslash\overleftarrow{r}]) \end{array}$$

As before a hook means 'in the previous state'. Then the first line of definition 3 tells us how to substitute for a variable, and the second line says that if P is unchanged then $\mathbb{B}(P)$ continues to hold. The $\mathbb{B}$ modality is about writes to shared variables, so register substitution sees through it.

Definition 4: Properties of the $\mathbb{B}$ modality

$$\begin{array}{rcll} \mathbb{B}(P) & \Rightarrow & P & \\ \mathbb{B}(P \wedge Q) & = & \mathbb{B}(P) \wedge \mathbb{B}(Q) & \\ \mathbb{B}(\mathbb{B}(P)) & = & \mathbb{B}(P) & \\ \mathbb{B}(P \vee Q) & \Leftarrow & \mathbb{B}(P) \vee \mathbb{B}(Q) & \\ \mathbb{B}(P) & \Leftarrow & P & \text{(if } P \text{ doesn't mention variables)} \\ \mathbb{B}(P \vee Q) & \Rightarrow & \mathbb{B}(P) \vee \mathbb{B}(Q) & \text{(if } P \text{ or } Q \text{ don't mention variables)} \\ (P \Rightarrow Q) & \Rightarrow & \mathbb{B}(P) \Rightarrow \mathbb{B}(Q) & \text{(if } P \Rightarrow Q \text{ is a tautology)} \end{array}$$

The first line of this definition allows a $\mathbb{B}$-modal assertion in an interference precondition to assert that P holds when the interference is received. These properties also apply, as we shall see, to the $\mathbb{U}$ and $\mathbb{S}$ofar modalities, because all three are, behind the scenes, versions of the 'since' modality of temporal logic and based on a $\forall$ quantification.

When all the occurrences of variables in an interference precondition are inside $\mathbb{B}$ modalities, as they are in fig. 9 and will be in all our examples until section 16, we can use a simple treatment of stability, equivalent to the SC stability rule 1.

Rule 3: External stability against $\mathbb{B}$-modal interference

Constraint assertion P is stable against external interference $Q \mid x := E$,
where every variable occurrence in Q is enclosed in a $\mathbb{B}$ modality, if
$\mathrm{sp}(P \wedge Q,\ x := E) \Rightarrow P$

We shall see in section 16 that this is a special case of the general external stability rule 7.

[6] Or might as well have witnessed . . . Once again program logic simplifies and abstracts

9.2 The same game on a new pitch

The checklist of table 4 is carefully worded to suit both SC and weak-memory concurrency proofs, in order to emphasise the continuity between earlier work and our own, but the switch from between-component assertions to constraint embroidery requires some adjustment. In weak memory check 1 applies to source postcondition and constraint embroidery, not entire preconditions as with SC, shifting the stability requirement to constraint embroidery. We have to allow for rule 2 wrapping bo constraint embroidery in a $\mathbb{B}$ modality and uo embroidery with $\mathbb{U}$. But despite all that, in principle it's the same checklist as before: check that preconditions derive from preceding postconditions, check that they are stable, check that the action of variable assignments is reflected in the guarantee.

We are playing the same rely/guarantee Owicki-Gries game that we played with SC, in a world where there are new kinds of relationship between components and, as we shall see, new forms of parallelism and corresponding new instabilities.

10 Checking the proof of weak-memory message passing

The lacing of fig. 8 is very tight. Commands must elaborate after the initial state is established, and in so order in each thread. Propagation may not be reordered. We seem to have reimposed much of the ordering that SC imposed. It looks plausible that the example will work in weak memory. Plausible isn't proof, but the proof is very easy, given the work we already did for the SC case and the fact that rule 3 lets us use the SC stability calculation.

In the sender of fig. 9, the check 1 tests are exactly as in the SC proof, even given the different treatment of the bo constraint and its $\mathbb{B}$ modality:

- on the $\alpha\rightarrow$a lo constraint $f{=}0 \Rightarrow$ true;
- on the a$\rightarrow$b bo constraint true $\wedge\, m{=}1 \Rightarrow m{=}1$ and $\mathbb{B}(m{=}1) \Rightarrow \mathbb{B}(m{=}1)$.

Stability in check 2 is immediate in each case because the receiver's guarantee is empty. Check 3 is passed because each variable assignment, with its precondition (the conjunction of its knot's embroideries) is transcribed to the guarantee.

In the receiver the implications for check 1 are just as in the SC proof, and there are no variable assignments so check 3 is immediate. Stability of the precondition of command c against true $\mid m := 1$ is exactly as (1) (same assertion, same interference, same calculation). Stability against $\mathbb{B}(m{=}1) \mid f := 1$ involves the $\mathbb{B}$ modality but, as before, we find that $m{=}1$ is stable:

$$\begin{array}{ll} & \mathrm{sp}((f{=}1 \Rightarrow m{=}1) \wedge \mathbb{B}(m{=}1), f := 1) \\ = & (\overleftarrow{f}{=}1 \Rightarrow m{=}1) \wedge \overleftarrow{\mathbb{B}(m{=}1)} \wedge m{=}1 \wedge f{=}1 \\ \Rightarrow & m{=}1 \\ \Rightarrow & f{=}1 \Rightarrow m{=}1 \end{array} \tag{3}$$

The remaining assertions are stable against the sender's interference because they don't mention f and because $m{=}1$ is stable. The checklist is completed; we have verified the proof. Note that in (3) the $\mathbb{B}$ modality has allowed us to conclude, just as in the SC proof, that when the receiver sees $f{=}1$ it will already have seen $m{=}1$.

10.1 (C) Compiling laced programs

The aim of our work is to provide rules that programmers can use to reliably fit parts of their program together. Providing proofs justifies confidence that the imposed constraints are effective. But we can use unembroidered lacing to describe an algorithm, we expect that programmers will do that, and we're sure that compilers would ignore embroidery even if we were to show it to them. So, given a laced program, how do we produce an implementation for a particular architecture?

In the examples in this paper we assume a wrapper which deals with initiation and termination of threads. We suppose that the wrapper sets up the initial state and propagates it to all threads before execution starts. We shall see in section 19 that this requires a uo constraint between the initial state and the commands of each thread: table 3 claims that MFENCE on Intel x86, sync on Power and dsb on ARM would be sufficient to provide it. Since uo is included in bo and lo, the initial constraints are then in place for each thread of our examples.

On x86 we could transcribe the commands of fig. 8 directly, since its constraints are implicit in the architecture. On Power, lwsync between the commands of the sender would provide a bo constraint, and an addr dependency would suffice to provide an lo constraint in the receiver. ARM is similar, using dsb rather than lwsync to provide uo order (which implies bo). The Power and ARM compilations would then provide exactly what (Pow, 2015) says is required to prohibit the execution pictured in fig. 3.

It is possible to imagine a compiler which could translate a laced program into a implementation program for a weak-memory target machine, though minimal barrier placement is a hard problem (Burckhardt et al., 2007; Huynh and Roychoudhury, 2007; Atig et al., 2011). It's also possible to imagine a program which could check, given an implementation program and a laced program, that the implementation provides the constraints of the lacing.

11 New instabilities

In the proof of weak-memory MP we considered the interference of one thread with another: that is, interference caused by *external* parallelism. We needed the $\mathbb{B}$ modality to constrain propagation, but otherwise the weak-memory proof was very similar to the SC proof. But weak memory introduces other kinds of parallelism:

Internal (lo) parallelism An assignment can be elaborated after the elaboration of the source and before the elaboration of the target of a constraint, interfering with the constraint's embroidery.

In-flight (bo) parallelism The writes made by variable assignments which are not bo ordered may be propagated out of so order, and then an so-later assignment will interfere with the interference precondition of an so-earlier.

Universal (uo) parallelism A variable assignment may undermine a thread-universal assertion in another thread. Dealt with in section 19.

Note that lo parallelism can apply to the embroidery of any constraint: it has to do with the lo ordering of the interfering assignment and the constraint's source and target. Also, bo parallelism is about the absence of bo constraints.

11.1 Internal (lo) parallelism

Consider thread 0 of fig. 10. Each of its assignments is laced directly to the initial assertion α using a bo constraint; we can use bo because the initial state is broadcast to all threads before any of them start

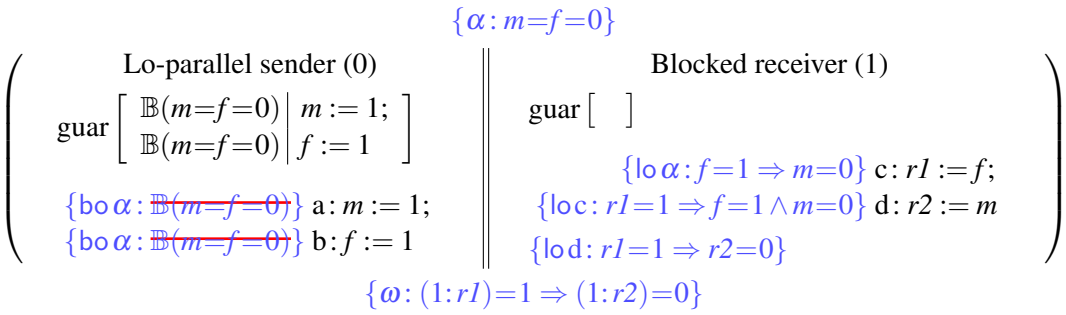
{α : m=f=0}
Lo-parallel sender (0)
guar [𝔹(m=f=0) | m := 1;]
[𝔹(m=f=0) | f := 1]
{bo α : 𝔹(m=f=0)} a : m := 1;
{bo α : 𝔹(m=f=0)} b : f := 1
Blocked receiver (1)
guar [ ]
{lo α : f=1 ⇒ m=0} c : r1 := f;
{lo c : r1=1 ⇒ f=1 ∧ m=0} d : r2 := m
{lo d : r1=1 ⇒ r2=0}
{ω : (1:r1)=1 ⇒ (1:r2)=0}

Figure 10: An unproof which ignores lo parallelism

executing. It appears that the embroidery is stable, just because there's no interference from the receiver, and the sender's guarantee is constructed as usual from the assignments and their preconditions. But note that this is an unproof, and we have struck through the embroidery in thread 0 because, as we shall see, it is unstable and therefore invalid.

Supposing for the moment that the embroidery were valid, we can imagine pairing this sender with the MP receiver of fig. 9. Then the precondition of the receiver's command c would be unstable against the second line of this sender's guarantee:

$$\begin{array}{ll} & \mathrm{sp}((f{=}1 \Rightarrow m{=}1) \wedge \mathbb{B}(m{=}f{=}0),\, f := 1) \\ = & ((\overleftarrow{f}{=}1 \Rightarrow m{=}1) \wedge \overleftarrow{\mathbb{B}(m{=}f{=}0)} \wedge m{=}\overleftarrow{f}{=}0 \wedge f{=}1 \\ \Rightarrow & f{=}1 \wedge m{=}0 \\ \not\Rightarrow & f{\neq}1 \vee m{=}1 \\ = & f{=}1 \Rightarrow m{=}1 \end{array} \tag{4}$$

So we can't prove successful message passing, which is as we should expect: experiment shows that the sender needs a bo constraint between a and b in order to succeed, and we don't have that in fig. 10.

But message passing is not the point of this example. The sender's guarantee is nonsense: it claims that each of the sender's writes $m{=}1$ and $f{=}1$ has to arrive at other threads in the initial state $m{=}f{=}0$, so that if either arrives first and changes the state, the other is blocked. Even weak memory isn't *that* silly.

Thread 1 of fig. 10 illustrates the silliness. The precondition of c is not the same as in the MP receiver of fig. 9: there we had $m{=}1$; here we have $m{=}0$. This precondition is stable under the stated interference of the lo-parallel sender:

$$\begin{array}{ll} & \mathrm{sp}((f{=}1 \Rightarrow m{=}0) \wedge \mathbb{B}(m{=}f{=}0),\, m := 1) \\ = & (f{=}1 \Rightarrow \overleftarrow{m}{=}0) \wedge \overleftarrow{\mathbb{B}(m{=}f{=}0)} \wedge \overleftarrow{m}{=}f{=}0 \wedge m{=}1 \\ \Rightarrow & f{=}0 \\ \Rightarrow & f{=}1 \Rightarrow m{=}0 \end{array} \tag{5}$$

$$\begin{array}{ll} & \mathrm{sp}((f{=}1 \Rightarrow m{=}0) \wedge \mathbb{B}(m{=}f{=}0),\, f := 1) \\ = & (\overleftarrow{f}{=}1 \Rightarrow m{=}0) \wedge \overleftarrow{\mathbb{B}(m{=}f{=}0)} \wedge m{=}\overleftarrow{f}{=}0 \wedge f{=}1 \\ \Rightarrow & m{=}0 \\ \Rightarrow & f{=}1 \Rightarrow m{=}0 \end{array} \tag{6}$$

The postcondition of c is $(f{=}1 \Rightarrow m{=}0) \wedge r1{=}f$, which implies $r1{=}1 \Rightarrow f{=}1 \wedge m{=}0$, the precondition of d.

That precondition is also stable against the sender's stated guarantee, each line of which requires $f{=}m{=}0$:

$$\begin{array}{ll} & \mathrm{sp}((r1{=}1 \Rightarrow f{=}1 \wedge m{=}0) \wedge \mathbb{B}(m{=}f{=}0),\ m := 1) \\ = & (r1{=}1 \Rightarrow f{=}1 \wedge \overleftarrow{m}{=}0) \wedge \overleftarrow{\mathbb{B}(m{=}f{=}0)} \wedge \overleftarrow{m}{=}f{=}0 \wedge m{=}1 \\ \Rightarrow & (r1{=}1 \Rightarrow f{=}1) \wedge f{=}0 \\ \Rightarrow & r1{\neq}1 \\ \Rightarrow & r1{=}1 \Rightarrow f{=}1 \wedge m{=}0 \end{array} \tag{7}$$

$$\begin{array}{ll} & \mathrm{sp}((r1{=}1 \Rightarrow f{=}1 \wedge m{=}0) \wedge \mathbb{B}(m{=}f{=}0),\ f := 1) \\ = & (r1{=}1 \Rightarrow \overleftarrow{f}{=}1 \wedge m{=}0) \wedge \overleftarrow{\mathbb{B}(m{=}f{=}0)} \wedge m{=}\overleftarrow{f}{=}0 \wedge f{=}1 \\ \Rightarrow & f{=}1 \wedge m{=}0 \\ \Rightarrow & r1{=}1 \Rightarrow f{=}1 \wedge m{=}0 \end{array} \tag{8}$$

The postcondition of d is $(r1{=}1 \Rightarrow f{=}1 \wedge m{=}0) \wedge r2{=}m$, which implies $r1{=}1 \Rightarrow r2{=}0$. So if command c reads $f{=}1$, command d *cannot* read $m{=}1$. This must be nonsense: the sender writes $m{=}1$, and the receiver can't be blocked from seeing it. Put it another way: (7) says that if the $m := 1$ interference arrives, then the $f := 1$ interference cannot have, and vice-versa with (8).

Our unproof overlooks weak-memory parallelism *within* a thread, which allows an assignment to interfere with the embroidery of a constraint on another command.

Definition 5: lo parallelism

An assignment a is *lo parallel* with a constraint b→c if tbere is an so path connecting a, b and c in which a is not constrained to come before b or after c.

We can use an SC stability calculation to check for lo stability, because there are no propagation effects inside a thread.

Rule 4: LO stability

Constraint embroidery P is LO stable against lo-parallel assignment A with precondition Q if

$$\mathrm{sp}(P \wedge Q,\ A) \Rightarrow P$$

The rule applies to all kinds of assignments: variable writes $x := E$, reads $r := x$ and calculations $r := E$. (But on some architectures definition 5 may not apply to reads and calculations: see the discussion of SCreg in section 17.2.)

Each of the assignments in the sender of fig. 10 is lo-parallel with the constraint of the other: a is so-between the ends of the α→b constraint, and b is so- but not lo-after the α→a constraint. In neither case do we have stability. We test the α→a embroidery against the interference of b:

$$\begin{array}{ll} & \mathrm{sp}(\mathbb{B}(m{=}0 \wedge f{=}0) \wedge \mathbb{B}(m{=}0 \wedge f{=}0),\ f := 1) \\ = & \mathrm{sp}(\mathbb{B}(m{=}0 \wedge f{=}0),\ f := 1) \\ = & \overleftarrow{\mathbb{B}(m{=}0 \wedge f{=}0)} \wedge m{=}0 \wedge \overleftarrow{f}{=}0 \wedge f{=}1 \\ \not\Rightarrow & \mathbb{B}(m{=}0 \wedge f{=}0) \end{array} \tag{9}$$

The case of the α→b embroidery and the interference of a is similar.

If we weaken the α→a embroidery to $\mathbb{B}(m{=}0)$ and the α→b embroidery to $\mathbb{B}(f{=}0)$, as in fig. 11, we have lo stability in the sender, since neither of those embroideries mentions the other assignment's variable. But now have external instability in the receiver. The α→c embroidery is unstable against both lines of the

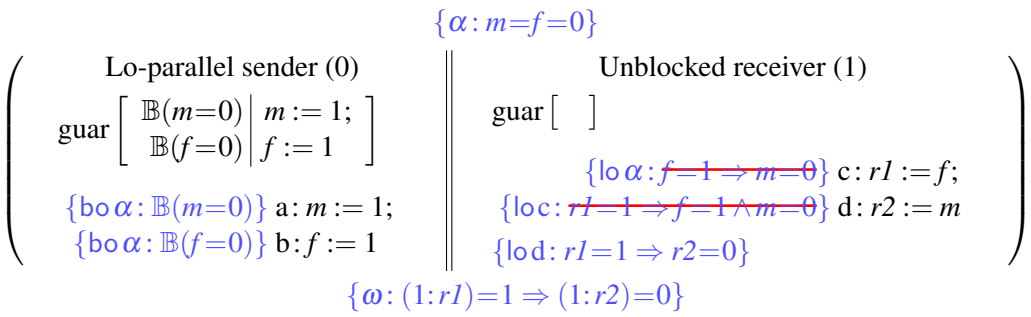
{α : m=f=0}
Lo-parallel sender (0)
guar [B(m=0) | m := 1;]
[B(f=0) | f := 1]
{bo α : B(m=0)} a : m := 1;
{bo α : B(f=0)} b : f := 1
Unblocked receiver (1)
guar [ ]
{lo α : f=1 ⇒ m=0} c : r1 := f;
{loc : r1=1 ⇒ f=1 ∧ m=0} d : r2 := m
{lo d : r1=1 ⇒ r2=0}
{ω : (1:r1)=1 ⇒ (1:r2)=0}

Figure 11: Lo parallelism and blocking removed, but still an unproof

sender's guarantee. Here's the $\mathbb{B}(f{=}0)\,|\,f := 1$ case:

$$\begin{array}{ll} & \mathrm{sp}((f{=}1 \Rightarrow m{=}0) \wedge \mathbb{B}(f{=}0), f := 1) \\ = & (\overleftarrow{f}{=}1 \Rightarrow m{=}0) \wedge \mathbb{B}(\overleftarrow{f}{=}0) \wedge \overleftarrow{f}{=}0 \wedge f{=}1 \\ \not\Rightarrow & f{\neq}1 \vee m{=}0 \end{array} \tag{10}$$

The c→d embroidery, which prevented the receiver from reading $m{=}1$ when $r1{=}1$, is unstable against $\mathbb{B}(m{=}0)\,|\,m := 1$:

$$\begin{array}{ll} & \mathrm{sp}((r1{=}1 \Rightarrow f{=}1 \wedge m{=}0) \wedge \mathbb{B}(m{=}0),\ m := 1) \\ = & (r1{=}1 \Rightarrow f{=}1 \wedge \overleftarrow{m}{=}0) \wedge \mathbb{B}(\overleftarrow{m}{=}0) \wedge \overleftarrow{m}{=}0 \wedge m{=}1 \\ \not\Rightarrow & r1{\neq}1 \vee f{=}1 \wedge m{=}0 \end{array} \tag{11}$$

Even though the silly blocking has gone, the sender's weakened guarantee still doesn't support reliable message-sending: the precondition of the MP receiver's c (see fig. 9) is unstable against the second line of the new sender's guarantee:

$$\begin{array}{ll} & \mathrm{sp}((f{=}1 \Rightarrow m{=}1) \wedge \mathbb{B}(f{=}0), f := 1) \\ = & ((\overleftarrow{f}{=}1 \Rightarrow m{=}1) \wedge \mathbb{B}(\overleftarrow{f}{=}0) \wedge \overleftarrow{f}{=}0 \wedge f{=}1 \\ \not\Rightarrow & f{\neq}1 \vee m{=}1 \\ = & f{=}1 \Rightarrow m{=}1 \end{array} \tag{12}$$

Note that in the MP sender of fig. 9 the assigments are ordered by bo and therefore by lo, so there is no lo parallelism. In the MP receiver of fig. 9 and the receiver of fig. 10 and fig. 11 the assignments are ordered by lo, so no lo parallelism there either.

11.2 In-flight (bo) parallelism

Consider thread 0 of fig. 12. We've eliminated lo parallelism by including an lo constraint between a and b. This time we do have stable embroidery, but we shall see that the interference precondition of the first line of the sender's guarantee, though derived transparently from the elaboration precondition of a, is invalid. If we were to trust that interference precondition, we would have another nonsense guarantee: the assertions in the receiver are stable against it, but if thread 2's c reads $f{=}1$ and d reads $m{=}0$ then, according to its precondition, e *cannot* read $m{=}1$.

The problem is propagation reordering. In thread 0 there is no bo constraint a→b to ensure that $m{=}1$ is propagated before $f{=}1$, and if they are propagated to thread 1 in the reverse order – $f{=}1$ before $m{=}1$ –

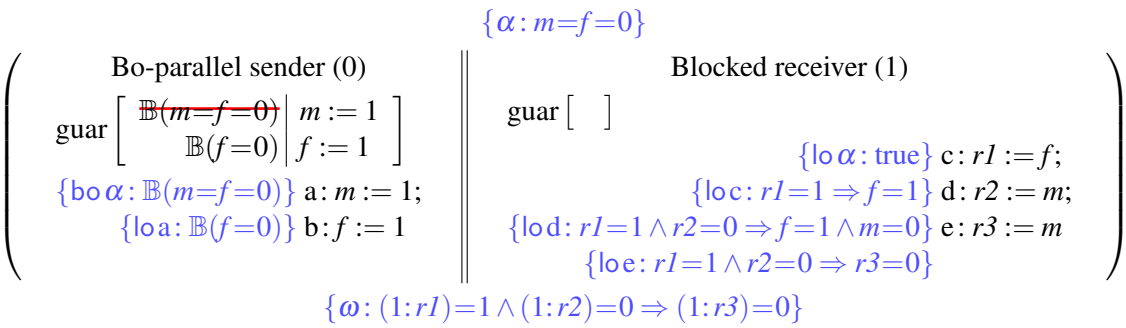
{α : m=f=0}
Bo-parallel sender (0)
guar [~~𝔹(m=f=0)~~ | m := 1 ; 𝔹(f=0) | f := 1]
{bo α : 𝔹(m=f=0)} a : m := 1;
{loa : 𝔹(f=0)} b : f := 1
Blocked receiver (1)
guar [ ]
{lo α : true} c : r1 := f;
{loc : r1=1 ⇒ f=1} d : r2 := m;
{lod : r1=1 ∧ r2=0 ⇒ f=1 ∧ m=0}
{loe : r1=1 ∧ r2=0 ⇒ r3=0}
e : r3 := m
{ω : (1:r1)=1 ∧ (1:r2)=0 ⇒ (1:r3)=0}

Figure 12: An unproof which ignores bo parallelism

then the precondition of the first line of the sender's guarantee, which insists that $f=0$, won't hold when $m=1$ arrives. In effect the interference generated by b is altering the interference generated by a. It's weak-memory *in-flight* or *bo-parallel* interference.

Definition 6: Bo **parallelism**

If variable assignments to distinct variables are so- but not bo-ordered then the so-later is *bo parallel* with the so-earlier and will interfere with its interference precondition.

Note the asymmetry of this definition: the second assignment is bo-parallel with the first, but not the first with the second. For justification, see below. The restriction to distinct variables is because writes to the same variable can't overtake each other: for explanation, see section 16.

As with all our examples until section 16, we once again consider only the special case of $\mathbb{B}$-modal interference preconditions. For the general rule, see section 16.

Rule 5: BO stability with $\mathbb{B}$-modal interference

Interference $P \mid x := E$ is BO stable against bo-parallel $Q \mid y := F$,
where both P and Q have no occurrences of variables outside a $\mathbb{B}$ modality, if
$\mathrm{sp}(P \wedge Q,\ x := E) \Rightarrow P$

Using this rule, we check stability of a's interference in fig. 12 against the interference of b. It's unstable:

$$\begin{array}{ll} & \mathrm{sp}(\mathbb{B}(m{=}0 \wedge f{=}0) \wedge \mathbb{B}(f{=}0),\ f := 1) \\ = & \mathrm{sp}(\mathbb{B}(m{=}0) \wedge \mathbb{B}(f{=}0),\ f := 1) \\ = & \overbrace{\mathbb{B}(m{=}0)} \wedge m{=}0 \wedge \overbrace{\mathbb{B}(f{=}0)} \wedge \overleftarrow{f} {=} 0 \wedge f{=}1 \\ \Rightarrow & \mathbb{B}(m{=}0) \wedge \overbrace{\mathbb{B}(f{=}0)} \wedge \overleftarrow{f} {=} 0 \wedge f{=}1 \\ \not\Rightarrow & \mathbb{B}(m{=}0) \wedge \mathbb{B}(f{=}0) \\ = & \mathbb{B}(m{=}0 \wedge f{=}0) \end{array} \tag{13}$$

We shouldn't weaken a's elaboration precondition, which is validly derived from the stable embroidery of its constraints, and gives us a postcondition which validates the elaboration precondition of b. But we must weaken a's interference precondition to regain in-flight stability.

In order that we can check for bo-parallel stability in our checklist, the laced program must declare what the interference generated by each assignment actually is. In this example there is a single line in the guarantee for each assignment, so it's quite easy to see what's going on, but things aren't always so simple.

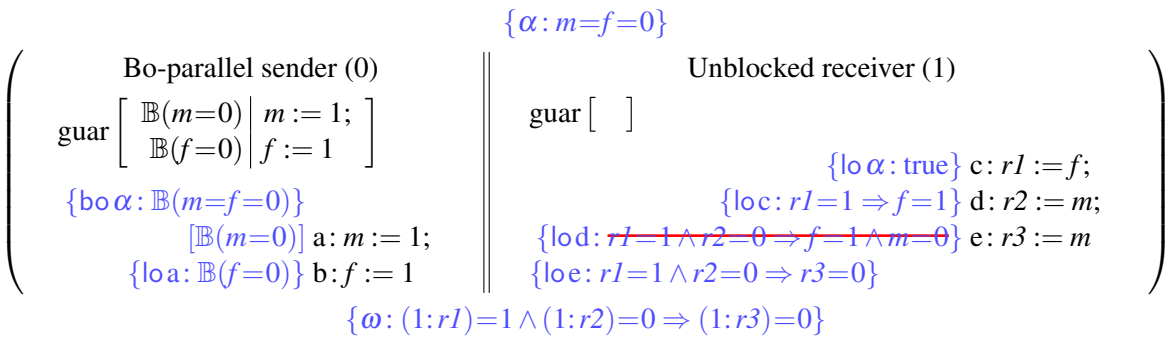
{α : m=f=0}
Bo-parallel sender (0)
guar [𝔹(m=0) | m := 1;]
[𝔹(f=0) | f := 1]
{bo α : 𝔹(m=f=0)}
[𝔹(m=0)] a : m := 1;
{loa : 𝔹(f=0)} b : f := 1
Unblocked receiver (1)
guar [ ]
{lo α : true} c : r1 := f;
{loc : r1=1 ⇒ f=1} d : r2 := m;
{lod : r1=1∧r2=0 ⇒ f=1∧m=0} e : r3 := m
{loe : r1=1∧r2=0 ⇒ r3=0}
{ω : (1:r1)=1∧(1:r2)=0 ⇒ (1:r3)=0}

Figure 13: Bo instability removed makes an unproof of silly blocking

Our solution is to append, after the knot which describes the elaboration precondition, a square-bracketed interference precondition. When it's absent, the elaboration precondition is the interference precondition. So we've declared, in fig. 13, that the interference precondition of a is $\mathbb{B}(m{=}0)$, and transcribed that to the sender's guarantee.

Fig. 13 removes the bo instability in the sender – a's precondition no longer mentions f, so is stable against the bo-parallel interference of b – and removes the silliness in the receiver – the precondition of e is now unstable against the first line of the sender's guarantee.

$$\begin{array}{ll} & \mathrm{sp}((r1{=}1\wedge r2{=}0\Rightarrow f{=}1\wedge m{=}0)\wedge\mathbb{B}(m{=}0),\ m:=1) \\ = & (r1{=}1\wedge r2{=}0\Rightarrow f{=}1\wedge \overleftarrow{m}{=}0)\wedge\overleftarrow{\mathbb{B}(m{=}0)}\wedge\overleftarrow{m}{=}0\wedge m{=}1 \\ \not\Rightarrow & r1{=}1\wedge r2{=}0\Rightarrow f{=}1\wedge m{=}0 \end{array} \tag{14}$$

We still don't get message passing between this sender and the MP receiver because of the instability demonstrated in (12).

Definition 6 and rule 5 deal with the interference of one variable assignment with the interference precondition of an earlier assignment. We don't have to check in the opposite direction (stability of later against earlier) because the elaboration precondition of the later assignment, and therefore its interference precondition, already take account of the earlier assignment. In fig. 13, for example, the elaboration precondition of command b cannot imply $\mathbb{B}(m{=}0)$ and therefore neither can its interference precondition.

11.3 Bo parallelism in the rely

Lo parallelism is entirely within a single thread, and bo parallelism can occur within a single thread, as we have shown. But when the interferences of more than one thread converge on another, we have to guard against in-flight interference between their separate guarantees. So when forming a rely for one thread from the guarantees of other threads, we have to form the bo-stable merge of those other threads' guarantees, weakening where necessary. Two examples of this will be seen in section 30; the rule we use will be revealed in section 16.

11.4 Still the same game

Check 2 of our checklist requires assertions to be stable against 'parallel assignments', which now includes both external and lo parallelism. Check 3 requires 'stable interference' in the guarantee; we have seen that

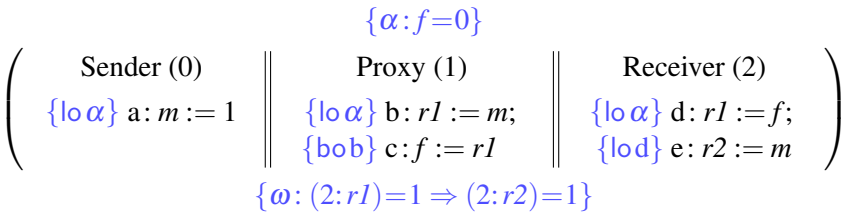

Figure 14: AlmostWRC: sending through a proxy

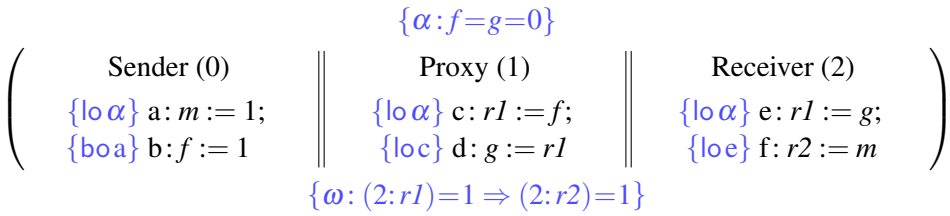

Figure 15: AlmostISA2: flagging through a proxy

we must check bo-parallel assignments for interference stability.

We are still playing the Owicki-Gries-Jones game: check that assertions are stable, check that the external effects of variable assignments are reflected in the guarantee.

12 Reads and the B modality

In the examples so far we have applied bo constraints between variable assignments. But because bo controls propagation of the writes underpinning the source postcondition, no matter what kind of component the source is, we can usefully use it in other contexts. Fig. 14 shows a variation of message-passing in which a proxy thread can receive the sender's message and pass it on to the receiver. Observe that the sender plus the proxy look rather like the sender of MP, extended across two threads – but they're not identical, because the proxy in fig. 14 might not read $m{=}1$.

We have seen in the MP example, in both SC and weak-memory treatments, that an interference precondition carries information about earlier writes in the same thread, so that we can assert an implication which is stable against that interference. The implication $f{=}1 \Rightarrow m{=}1$ in the MP proofs works off the interference description $m{=}1 \mid f{=}1$ in the SC case and $\mathbb{B}(m{=}1) \mid f{=}1$ in the weak memory case (using the fact that $\mathbb{B}(m{=}1) \Rightarrow m{=}1$ when the interference hits). It's possible in the same way to use $\mathbb{B}$-modal assertions, effectively extending bo between threads. Fig. 15 shows a variation of an example which is considered rather subtle in discussions of the Power architecture. The sender is as in MPlaced; the proxy thread reads the sender's flag f and notifies the receiver via a proxy flag g. The subtlety is that there is no need for a bo constraint between commands c and d in the proxy: lo everywhere but in the sender is enough. This time it is the MP receiver that seems to be extended across two threads.

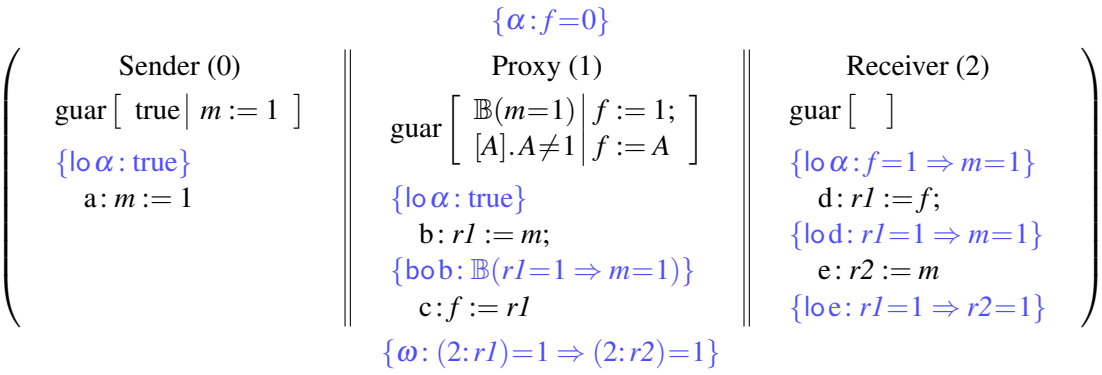
{α : f=0}
Sender (0)
guar [true | m := 1]
{lo α : true}
a : m := 1
Proxy (1)
guar [𝔹(m=1) | f := 1;]
[A].A≠1 | f := A
{lo α : true}
b : r1 := m;
{bo b : 𝔹(r1=1 ⇒ m=1)}
c : f := r1
{ω : (2:r1)=1 ⇒ (2:r2)=1}
Receiver (2)
guar [ ]
{lo α : f=1 ⇒ m=1}
d : r1 := f;
{lo d : r1=1 ⇒ m=1}
e : r2 := m
{lo e : r1=1 ⇒ r2=1}

Figure 16: A proof of almostWRC

12.1 Bo from read to write

Fig. 16 is a proof of almostWRC. The rely of the sender – the interference it can experience – is the bo-stable union of the guarantees of the proxy and the receiver: because the receiver's guarantee is empty, that means just the guarantee of the proxy. The precondition of command a is implied by the initial assertion α (check 1); it's stable because it's a constant (check 2); a's interference is included in the guarantee (check 3).

The rely of the proxy is in effect the guarantee of the sender. The precondition of b is implied by the initial assertion α (check 1), it's stable because it's a constant (check 2); b isn't a variable assignment, so there's nothing for the guarantee (check 3).

The postcondition of b is true $\wedge\, r1{=}m$, which implies $r1{=}1 \Rightarrow m{=}1$, which we wrap in $\mathbb{B}$ to embroider the bo b$\rightarrow$c constraint (check 1). That embroidery is stable against the interference of the sender, just because $m{=}1$ is stable (check 2):

$$\begin{array}{rl} & \overrightarrow{\mathrm{sp}}(\mathbb{B}(r1{=}1 \Rightarrow m{=}1),\ m := 1) \\ = & \overrightarrow{\mathbb{B}(r1{=}1 \Rightarrow m{=}1)} \wedge (r1{=}1 \Rightarrow \overleftarrow{m}{=}1) \wedge m{=}1 \\ \Rightarrow & \overrightarrow{\mathbb{B}(r1{=}1 \Rightarrow m{=}1)} \wedge m{=}1 \\ \Rightarrow & \overrightarrow{\mathbb{B}(r1{=}1 \Rightarrow m{=}1)} \wedge (r1{=}1 \Rightarrow m{=}1) \\ \Rightarrow & \mathbb{B}(r1{=}1 \Rightarrow m{=}1) \end{array} \tag{15}$$

A novelty is the treatment of the interference generated by command c. Because precondition and assignment mention register $r1$, which by definition is local to the thread, we must quotient the register, giving the interference description

$$[r1].\,\mathbb{B}(r1{=}1 \Rightarrow m{=}1)\ |\ f := r1 \tag{16}$$

To check stability against quotiented interference, we must consider all the possible values of the quotiented registers:

Rule 6: Stability against quotiented interference

P is stable against $[\overline{ns}].\,Q \mid x := E$ if it is stable against $Q[\overline{ns}\backslash\overline{fs}] \mid x := E[\overline{ns}\backslash\overline{fs}]$ where $\overline{fs}$ are fresh names. Similarly for internal interference $[\overline{ns}].\,Q \mid r := E$ and $[\overline{ns}].\,Q \mid r := x$.

We can split (16) for the two cases $r1{=}1$ (17) and $r1{\neq}1$ (18):

$$\mathbb{B}(m{=}1) \mid f := 1 \tag{17}$$
$$[r1].\, r1{\neq}1 \mid f := r1 \tag{18}$$

(17) is the first line of the proxy's guarantee. (18) is the second line, renaming $r1$ to A. That completes check 3, since there is no other assignment to be bo-parallel with.

The rely of the receiver is the bo-stable union of the guarantees of sender and proxy. In this example we can simply merge them: the interferences of the proxy are immune to the interferences of the sender (in line 1, $m{=}1$ is stable, and in line 2, the precondition doesn't mention m or f); the interference preconditions of the sender don't mention f, which is the only variable the proxy writes. We have the same receiver with the same guarantee and embroidery as in the MP proof (fig. 9), but with a new rely we'll have to make some new checks.

We already showed, when checking fig. 9, that the receiver's embroidery is stable against the sender's guarantee and the first line of the proxy's guarantee. The second line of the proxy's guarantee is new, but the precondition of d contains the only assertion that mentions f. Stability holds because the interference establishes $f{\neq}1$:

$$\begin{array}{ll} & \mathrm{sp}((f{=}1 \Rightarrow m{=}1) \wedge A{\neq}1,\, f := A) \\ = & (\overleftarrow{f}{=}1 \Rightarrow m{=}1) \wedge A{\neq}1 \wedge f{=}A \\ \Rightarrow & f{\neq}1 \\ \Rightarrow & f{=}1 \Rightarrow m{=}1 \end{array} \tag{19}$$

With that we've completed verification of the proof of almostWRC. Causally, the assignment $r1 := m$ is bo-before $f := r1$; if the proxy reads from $m := 1$ so that $r1{=}1$, then $m := 1$ is also bo-before $f := 1$, just as in fig. 9. Our proof agrees.

12.2 The B modality transmitted by interference

Fig. 17 shows a proof of almostISA2. The sender is exactly as in the MP proof, and the proxy's guarantee doesn't interfere with any variables mentioned in the sender's assertions, so the earlier verification stands.

The proof-subtlety is in the precondition of command c: $f{=}1 \Rightarrow \mathbb{B}(m{=}1)$ echoes $f{=}1 \Rightarrow m{=}1$ in the proofs of MP. It's trivially true in initial state α because $f{=}0$, and is stable under the interference of the sender:

$$\begin{array}{ll} & \mathrm{sp}((f{=}1 \Rightarrow \mathbb{B}(m{=}1)) \wedge \mathrm{true}, m := 1) \\ = & (f{=}1 \Rightarrow \overleftarrow{\mathbb{B}(m{=}1)} \wedge \overleftarrow{m}{=}1) \wedge m{=}1 \\ = & (f{=}1 \wedge \overleftarrow{\mathbb{B}(m{=}1)} \wedge \overleftarrow{m}{=}1 \wedge m{=}1) \vee f{\neq}1 \wedge m{=}1 \\ \Rightarrow & \mathbb{B}(m{=}1) \vee f{\neq}1 \\ = & f{=}1 \Rightarrow \mathbb{B}(m{=}1) \end{array} \tag{20}$$

$$\begin{array}{ll} & \mathrm{sp}((f{=}1 \Rightarrow \mathbb{B}(m{=}1)) \wedge \mathbb{B}(m{=}1),\, f := 1) \\ = & (\overleftarrow{f}{=}1 \Rightarrow \mathbb{B}(m{=}1)) \wedge \mathbb{B}(m{=}1) \wedge f{=}1 \\ \Rightarrow & \mathbb{B}(m{=}1) \\ \Rightarrow & f{=}1 \Rightarrow \mathbb{B}(m{=}1) \end{array} \tag{21}$$

The precondition of d is stable, because $m{=}1$ is stable. The interference is

$$[r1].\, r1{=}1 \Rightarrow \mathbb{B}(m{=}1) \mid g := r1 \tag{22}$$

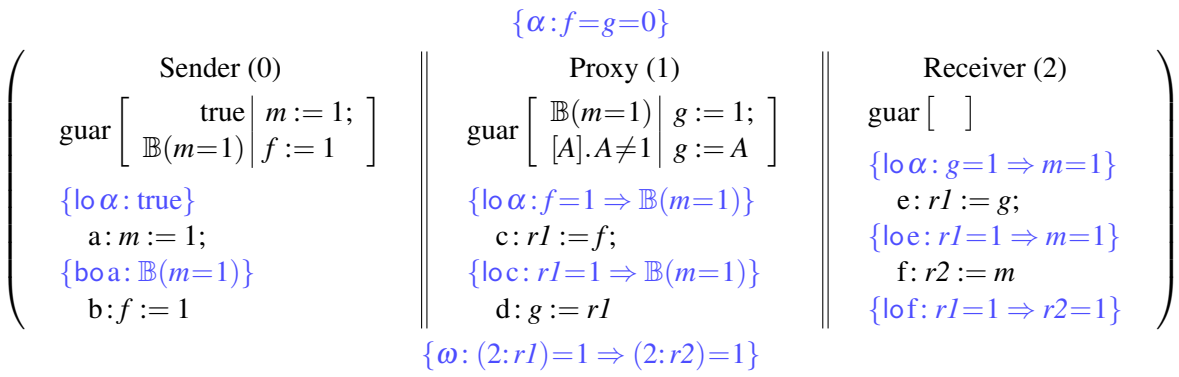
{α : f=g=0}
Sender (0)
guar [true | m := 1; ; 𝔹(m=1) | f := 1]
{lo α : true}
a : m := 1;
{bo a : 𝔹(m=1)}
b : f := 1
Proxy (1)
guar [𝔹(m=1) | g := 1; ; [A].A≠1 | g := A]
{lo α : f=1 ⇒ 𝔹(m=1)}
c : r1 := f;
{lo c : r1=1 ⇒ 𝔹(m=1)}
d : g := r1
{ω : (2:r1)=1 ⇒ (2:r2)=1}
Receiver (2)
guar [ ]
{lo α : g=1 ⇒ m=1}
e : r1 := g;
{lo e : r1=1 ⇒ m=1}
f : r2 := m
{lo f : r1=1 ⇒ r2=1}

Figure 17: A proof of almostISA2

which we can analyse just as we did (16), to give two entries in the proxy's guarantee.

The proof of the receiver is as in almostWRC, reading g for f. Its rely is effectively the same as in almostWRC, with the same reading, since the second line of the sender's guarantee writes to f, which isn't used by the receiver. So our earlier verification stands.

Causally, if $r1 := f$ reads from $f := 1$, which is bo-after m:=1, then it, too, is bo-after $m := 1$, and then so is $g := r1$. Bo extends from sender to proxy, and the $\mathbb{B}(m{=}1)$ modality in the proof tracks the ordering.

13 Auxiliary information in composite writes

Neither of the examples almostWRC and almostISA2 is quite the litmus test on which it is based. In each case the proxy reads a value and then writes it out again: if the value it reads is 1, it must write 1. In the WRC and ISA2 litmus tests of (Pow, 2015) the proxy always writes 1, and the proof challenge is to show that everything still works. Fig. 18, then, is the real WRC litmus test. It's laced exactly like almostWRC, but it has a slightly more complicated postcondition because we're not interested in the case in which the proxy reads 0 from m and the receiver reads 1 from f.

To make a proof we need somehow to distinguish the write-1-after-read-0 and write-1-after-read-1 interferences of the proxy. As usual in the Owick-Gries tradition, we resort to auxiliary assignments which record extra information to help us make a proof. It's necessary that the execution of the program, so far as regular variables and registers are concerned, is unaffected when auxiliary assignments are deleted. In SC proofs we mustn't assign the value of an auxiliary variable to a regular register, and we can't use auxiliary registers in conditional expressions. In weak memory proofs we must also control use of constraints between auxiliary assignments and regular components, which disappear when auxiliary assignments are deleted. The details of our treatment of constraints are given in section 31, but we needn't consider them in this example because it doesn't include any purely-auxiliary assignments.

In fig. 19 we've added auxiliary information to variable f – so-called 'auxiliary extension'. Command c writes 1, as in fig. 18, but also the value that was read by command b. That is, it writes either $(1,0)$ or $(1,1)$: the first part is the regular part of the write, the second part an auxiliary value. We imagine that the extended value is propagated as a unit. Command d in the receiver reads the regular part into $r1$ and ignores the auxiliary part.

Note that in the initial state assertion we've had to change $f{=}0$ into $f{=}(0,0)$. But note also that the auxiliary extension to f is completely irrelevant to execution, so much so that the receiver doesn't even look at it. This

$$\{\alpha : f=0\}$$

$$\left(\begin{array}{c|c|c} \text{Sender (0)} & \text{Proxy (1)} & \text{Receiver (2)} \\ \{\mathsf{lo}\,\alpha\}\ \mathrm{a}: m := 1 & \{\mathsf{lo}\,\alpha\}\ \mathrm{b}: r1 := m; & \{\mathsf{lo}\,\alpha\}\ \mathrm{d}: r1 := f; \\ & \{\mathsf{bob}\}\ \mathrm{c}: f := 1 & \{\mathsf{lo}\,\mathsf{d}\}\ \mathrm{e}: r2 := m \end{array}\right)$$

$$\{\omega : (1{:}r1){=}1 \wedge (2{:}r1){=}1 \Rightarrow (2{:}r2){=}1\}$$

Figure 18: WRC: sending through an inattentive proxy

$$\{\alpha : f=(0,0)\}$$

$$\left(\begin{array}{c|c|c} \text{Sender (0)} & \text{Proxy (1)} & \text{Receiver (2)} \\ \{\mathsf{lo}\,\alpha\}\ \mathrm{a}: m := 1 & \{\mathsf{lo}\,\alpha\}\ \mathrm{b}: r1 := m; & \{\mathsf{lo}\,\alpha\}\ \mathrm{d}: r1, \_ := f; \\ & \{\mathsf{bob}\}\ \mathrm{c}: f := 1, r1 & \{\mathsf{lo}\,\mathsf{d}\}\ \mathrm{e}: r2 := m \end{array}\right)$$

$$\{\omega : (1{:}r1){=}1 \wedge (2{:}r1){=}1 \Rightarrow (2{:}r2){=}1\}$$

Figure 19: WRC with auxiliary extension to variable f

is as it should be: auxiliary information is purely to facilitate proof.

Figure 20 shows our proof. We stacked the threads to fit on the page.

The sender's rely is the bo-stable union of the guarantees of the proxy and the receiver, which is easy to make since the receiver's guarantee is empty. The rely doesn't matter anyway, because the sender's only assertion is the constant true. The checks are just those we made for the almostWRC sender.

The proxy's rely is the union of the guarantee of the sender and the empty guarantee of the receiver – again, no inter-thread bo parallelism to detain us. The postcondition of b is $f=(0,0) \wedge r1{=}m$, which implies $f=(0,0) \wedge (r1{=}1 \Rightarrow m{=}1)$, which we wrap in a $\mathbb{B}$ modality to embroider the bo b$\rightarrow$c constraint. That embroidery is stable, since $f=(0,0)$ and $m{=}1$ are each stable under the sender's interference. The interference of c is therefore

$$[r1].\mathbb{B}(f=(0,0) \wedge (r1{=}1 \Rightarrow m{=}1)) \,|\, f := 1, r1 \tag{23}$$

The first line of the guarantee is the case when $r1{=}1$, the second when $r1{\neq}1$. The strongest postcondition of c implies $f=(1,r1)$ and thus the thread postcondition, which is stable because neither sender nor receiver writes to f.

The rely of the receiver is the bo-stable union of the guarantees of the sender and the proxy. There is no effect of inter-thread bo parallelism: the sender's interference doesn't depend on f, and the proxy's interference is stable against $m := 1$. The precondition of d is very familiar, reading $f=(1,1)$ for $f{=}1$, and is stable for familiar reasons. The strongest postcondition of d is

$$(f=(1,1) \Rightarrow m{=}1) \wedge r1, \_{=}f \tag{24}$$

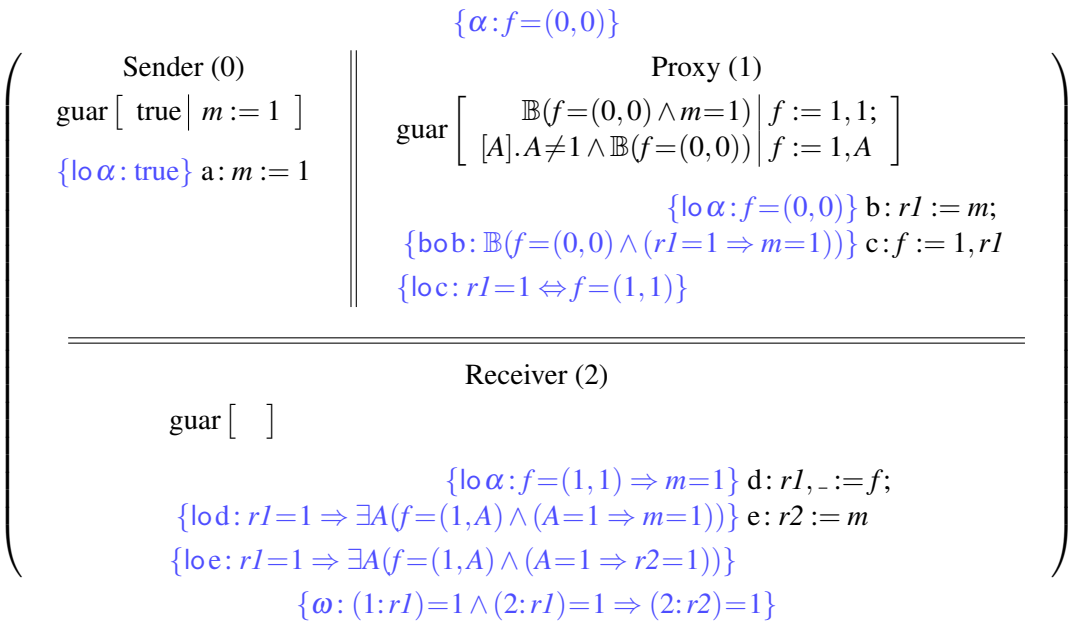
{α : f=(0,0)}
Sender (0)
guar [true | m := 1]
{lo α : true} a : m := 1
Proxy (1)
guar [𝔹(f=(0,0) ∧ m=1) | f := 1,1;
[A].A≠1 ∧ 𝔹(f=(0,0)) | f := 1,A]
{lo α : f=(0,0)} b : r1 := m;
{bob : 𝔹(f=(0,0) ∧ (r1=1 ⇒ m=1))} c : f := 1, r1
{loc : r1=1 ⇔ f=(1,1)}
Receiver (2)
guar [ ]
{lo α : f=(1,1) ⇒ m=1} d : r1, _ := f;
{lod : r1=1 ⇒ ∃A(f=(1,A) ∧ (A=1 ⇒ m=1))} e : r2 := m
{loe : r1=1 ⇒ ∃A(f=(1,A) ∧ (A=1 ⇒ r2=1))}
{ω : (1:r1)=1 ∧ (2:r1)=1 ⇒ (2:r2)=1}

Figure 20: A proof of WRC with auxiliary extension

which implies the embroidery of d→e. That embroidery is stable against each component of the rely:

$$\begin{array}{ll} & \mathrm{sp}((r1{=}1 \Rightarrow \exists A(f{=}(1,A) \wedge (A{=}1 \Rightarrow m{=}1))) \wedge \mathrm{true},\ m := 1) \\ = & (r1{=}1 \Rightarrow \exists A(f{=}(1,A) \wedge (A{=}1 \Rightarrow \overleftarrow{m}{=}1))) \wedge m{=}1 \\ \Rightarrow & (r1{=}1 \Rightarrow \exists A(f{=}(1,A))) \wedge m{=}1 \\ \Rightarrow & r1{=}1 \Rightarrow \exists A(f{=}(1,A) \wedge m{=}1) \\ \Rightarrow & r1{=}1 \Rightarrow \exists A(f{=}(1,A) \wedge (A{=}1 \Rightarrow m{=}1)) \end{array} \tag{25}$$

$$\begin{array}{ll} & \mathrm{sp}((r1{=}1 \Rightarrow \exists A(f{=}(1,A) \wedge (A{=}1 \Rightarrow m{=}1))) \wedge \mathbb{B}(f{=}(0,0) \wedge m{=}1),\ f := 1,1) \\ = & \left(\begin{array}{l} (r1{=}1 \Rightarrow \exists A(\overleftarrow{f}{=}(1,A) \wedge (A{=}1 \Rightarrow m{=}1))) \\ \wedge\ \overleftarrow{\mathbb{B}(f{=}(0,0) \wedge m{=}1)} \wedge \overleftarrow{f}{=}(0,0) \wedge m{=}1 \wedge f{=}(1,1) \end{array} \right) \\ \Rightarrow & f{=}(1,1) \wedge m{=}1 \\ \Rightarrow & \exists A(f{=}(1,A) \wedge (A{=}1 \Rightarrow m{=}1)) \\ \Rightarrow & r1{=}1 \Rightarrow \exists A(f{=}(1,A) \wedge (A{=}1 \Rightarrow m{=}1)) \end{array} \tag{26}$$

$$\begin{array}{ll} & \mathrm{sp}((r1{=}1 \Rightarrow \exists A(f{=}(1,A) \wedge (A{=}1 \Rightarrow m{=}1))) \wedge \mathbb{B}(f{=}(0,0) \wedge A{\neq}1),\ f := 1,A) \\ = & \left(\begin{array}{l} (r1{=}1 \Rightarrow \exists A(\overleftarrow{f}{=}(1,A) \wedge (A{=}1 \Rightarrow m{=}1))) \\ \wedge\ \overleftarrow{\mathbb{B}(f{=}(0,0) \wedge A{\neq}1)} \wedge \overleftarrow{f}{=}(0,0) \wedge A{\neq}1 \wedge f{=}(1,A) \end{array} \right) \\ \Rightarrow & A{\neq}1 \wedge f{=}(1,A) \\ \Rightarrow & \exists A(f{=}(1,A) \wedge A{\neq}1) \\ \Rightarrow & \exists A(f{=}(1,A) \wedge (A{=}1 \Rightarrow m{=}1)) \\ \Rightarrow & r1{=}1 \Rightarrow \exists A(f{=}(1,A) \wedge (A{=}1 \Rightarrow m{=}1)) \end{array} \tag{27}$$

The strongest postcondition of e is $r1{=}1 \Rightarrow \exists A(f{=}(1,A) \wedge (A{=}1 \Rightarrow m{=}1)) \wedge r2{=}m$, which implies the e→postcondition embroidery; that embroidery is stable because $\exists A(f{=}(1,A))$ and $r2{=}1$ are stable. The program postcondition ω is implied by the conjunction of the final assertions of the proxy and the receiver

and, as usual, true from the sender.

It's all rather neat. The implication $f=(1,1) \Rightarrow m=1$ is implied in the proxy's interference, and the information $r1=1 \Leftrightarrow f=(1,1)$ in its postcondition. The stability of assertions in the receiver comes, more or less, from the stability of $m=1$ and $f=(1,1)$.

14 The lo, bo fragment

Although we've concentrated on simple, almost trivial examples, we've done enough to show how elaboration-order lo and before-order bo can be used to impose orderings on a weak-memory program, even in what are thought to be subtle litmus tests. Lo prevents elaboration reordering; bo prevents propagation reordering. Embroidered with assertions they can be used, according to well-established and well-understood reasoning principles, to show that a program has enough ordering to work.

Well-established and well-understood principles they may be, but weak-memory concurrency proofs aren't simple or trivial. Even in SC it's hard to invent algorithms, let alone prove that they work. Weak memory is worse, as you will surely have observed.

Most of our examples have involved a great deal of lo ordering, almost restoring SC in most cases. We haven't added as much as perhaps you might think, because some of that ordering will be implicit in particular weak-memory architectures: see section 16. And our examples so far are small, often with nothing going on but some weak memory effect. Counteracting that effect allows little room for internal parallelism. We shall, however, see examples of internal parallelism in section 30.

15 Structured commands and disjunction of knots

Reordering of elaboration and/or propagation introduces uncertainty about program execution. Constraints remove unhelpful uncertainty. Structured commands – conditionals and loops – allow alternative so-tree paths and introduce helpful uncertainty. But they also introduce complexity: conditionals allow instances of the same component on different paths, and loops allow multiple instances of the same component on the same path.

Hoare logic for SC gives structured commands pre and postconditions. We rely entirely on constraints that relate primitive components – assignments, asserts and control expressions. As usual we introduce our treatment by considering examples.

15.1 Conditionals

A conditional 'if E then $C1$ else $C2$ fi' introduces alternative so-tree paths: one path includes $C1$, the other $C2$. An SC rule can use disjunction of postconditions to deal with the problem

$$\frac{\{P \wedge E\}\, C1\, \{Q1\} \quad \{P \wedge \neg E\}\, C2\, \{Q2\}}{\{P\}\ \mathsf{if}\ E\ \mathsf{then}\ C1\ \mathsf{else}\ C2\ \mathsf{fi}\ \{Q1 \vee Q2\}}\ \text{SC conditional} \tag{28}$$

Our logic doesn't give pre and postconditions to conditionals so we can't use this technique. Instead we use disjunction of knots.

Our example is a 'conditional receiver' for MP which reads f, and then reads m only if f was 1. Unlabelled and unlaced we might write

$$r1 := f;\ \mathsf{if}\ r1=1\ \mathsf{then}\ r2 := m\ \mathsf{fi} \tag{29}$$

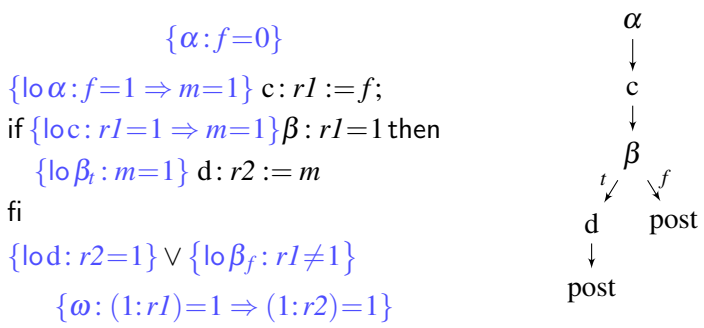

Figure 21: A conditional MP receiver and its so tree

Labelled, laced and embroidered we have the thread in fig. 21. In the same figure we show the thread's so tree. We expand the conditional into a tree with two arms, and append post, which stands for the thread postcondition, to each arm. We label the arrows exiting control expression β with t and f to indicate that one path is to be taken when β evaluates to true, the other when it is false. Note that the tree is derived entirely syntactically: it would have two arms even if β had a constant value,

Command c of fig. 21 is laced after the initial assertion. We lace control expression β after c so that it can test what c read, and we lace d after β_t so that it is only elaborated when c read 1. If execution follows the path including β_t then, this proof claims, d will certainly read 1. If execution follows the other path then certainly $r1 \neq 1$. The thread postcondition is a disjunction of knots, one covering the '$\beta_t \rightarrow$d$\rightarrow$post' path in the so tree, the other the '$\beta_f \rightarrow$post' path.

15.2 The constraint-coverage principle

A thread program contains a single instance of each command and each control expression. The so tree of the thread may contain multiple instances of a component; the program, in effect, conflates all those instances into one. In fig. 21, for example, the thread-postcondition component post appears twice. Constraints apply between instances of their source and target; to conflate those instances and their constraints we must apply the disjunction of the instance-constraints to the program component.

If a component were constrained on some paths but not all, the overall constraint would be true – i.e. no constraint. That would be useless as well as confusing.

Definition 7: Constraint-coverage principle

If a component is constrained at all, then for all paths in the so tree to an instance of that component, there must be a disjunct in the component's knot which depends on one of the nodes of that path.

The program in fig. 21 obeys the constraint-coverage principle. The sources of the constraints on c and β appear on every path to those components; the source of the constraint on d appears on every path to d; there's a disjunct in the thread-postcondition knot, which constrains termination, for each of the paths through the conditional.

15.3 Proof of the conditional receiver

The proof is in fig. 22. The sender is as in fig. 9. The precondition of command c in the receiver is implied by initial assertion α, as before in fig. 9; it is stable against the sender's guarantee, as before; there is no entry in the guarantee, as before. The precondition $r1=1 \Rightarrow m=1$ of control expression β is implied by the

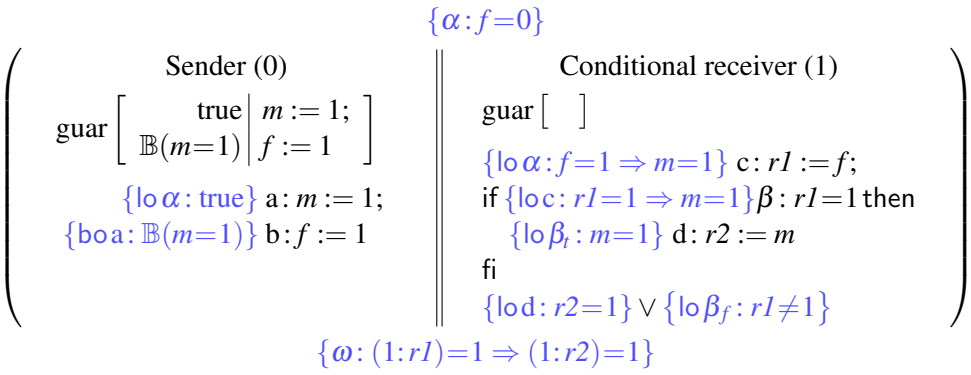
{α : f=0}
Sender (0)
guar [true | m := 1; ; 𝔹(m=1) | f := 1]
{lo α : true} a : m := 1;
{bo a : 𝔹(m=1)} b : f := 1
Conditional receiver (1)
guar [ ]
{lo α : f=1 ⇒ m=1} c : r1 := f;
if {lo c : r1=1 ⇒ m=1} β : r1=1 then
{lo β_t : m=1} d : r2 := m
fi
{lo d : r2=1} ∨ {lo β_f : r1≠1}
{ω : (1:r1)=1 ⇒ (1:r2)=1}

Figure 22: Proof of conditional message-passing

postcondition $(f=1 \Rightarrow m=1) \wedge r1=m$ of c, as with the precondition of d in fig. 9; it's stable, as before.

Definition 8: Postcondition of a control expression

A control expression $\gamma : E$ with elaboration precondition P has two postconditions, one for each of its outcomes. The γ_t outcome has $P \wedge E$; the γ_f outcome has $P \wedge \neg E$.

So for the $\beta_t \rightarrow$d constraint, the source postcondition is $(r1=1 \Rightarrow m=1) \wedge r1=1$, which implies the $m=1$ embroidery. That embroidery is stable against the rely. The postcondition of d implies $r2=1$, which is the embroidery on the d$\rightarrow$post constraint; once again it is stable.

For the $\beta_f \rightarrow$post constraint the source postcondition is $(r1=1 \Rightarrow m=1) \wedge r1 \neq 1$, which implies $r1 \neq 1$, which is stable against the rely. The disjunction $r2=1 \vee r1 \neq 1$ of the conjunctions of the assertions in each knot gives us the final postcondition that we require.

15.4 Conditional lacing summary

Lacing is already restricted so that it reinforces so. The constraint-coverage principle further restricts it so that if a knot outside a conditional appeals to one of its arms, then it must be disjoined with a knot that appeals to the other arm. In lacing basic commands and control expressions our intention is to allow the programmer to express the minimal constraints that make the program work. Conditionals (and loops) don't exist at the machine level: there are only instructions and jumps. Focusing on commands and control expressions allows us to work in the same sort of way.

15.5 (C) Compiling constraints with a control-expression source

As programmers we believe that barriers, which divide one part of an execution from another, are expensive, but that command orderings, which put one command before another, are cheap. Our lo constraint corresponds to a command-ordering dependency, which makes it cheaper than bo, which needs a barrier, which in turn is cheaper than uo, which needs a very big barrier indeed.[7]

On two important architectures, Power and ARM, lo constraints whose source is a control expression must be implemented with a branch plus a barrier (isync on Power, isb on ARM). It's the most lightweight of their

[7] Hardware architects encourage our belief that barriers are costly, and don't discourage our beliefs about their relative costs, but are tight-lipped about the magnitudes, so we're somewhat guessing.

$\{\mathsf{lo}\,\alpha\}\mathrm{c}: r1 := f;$
if $\{\mathsf{lo}\,\mathrm{c}\}\beta : r1 = 1$ then
  $\{\mathsf{lo}\,\mathrm{c}; \mathsf{lo}\,\beta_t\}\mathrm{d}: r2 := m$
fi
$\{\mathsf{lo}\,\mathrm{d}\} \vee \{\mathsf{lo}\,\beta_f\}$

Figure 23: Alternative, more execution-efficient, lacing of the conditional MP receiver

barriers, but to a programmer it's annoying. Surely it shouldn't be necessary to use heavy machinery to constrain something as simple as a conditional read. We believe that the barrier slows execution considerably[8] so we'd like to avoid it if we can. In many cases, luckily, we can do so very easily.

In the MP receiver of fig. 9, $\mathrm{d}: r2 := m$ is lo-constrained to be after $\mathrm{c}: r1 := f$. In fig. 21 the same is true, but the constraint is indirect, through the control expression $\beta : r1 = 1$. If we were to lace the thread slightly differently, making d depend on whatever β depends on, as in fig. 23, then we can argue that a barrier is unnecessary. It is irrelevant whether β is elaborated before, after or at the same time as d: if d is speculatively elaborated then, since it is on the same so path as β_t and lo-follows the same command, successful speculation can only take place in a state in which β either has delivered or will deliver true. The $\beta_t \rightarrow \mathrm{d}$ constraint is then operationally redundant and need not be implemented, though its embroidery may be safely used in the proof.

The isync/isb barrier is expensive, we believe, because it forces the hardware to order all so-before-barrier reads before all so-after-barrier reads. Given explicit lacing, we can see which reads need to be ordered, and we can use explicit lo constraints between them. A clever compiler might be able to notice that implementation of fig. 21 only needs an lo constraint $\mathrm{c} \rightarrow \mathrm{d}$. Expert programmers will know the trick, and will expect a compiler to implement it. But even if a compiler inserts a barrier when it doesn't really need to, we shall still get the result we require.

Although it can easily be eliminated in simple examples, isync is by no means redundant. In larger and more intricate examples it can be necessary to reserve a register for each multiply-constrained read in order to compile its lo dependencies. At some point pressure on registers would become overwhelming, and then it would be best to use ctrl-isync.

One of our intentions in developing a program logic (at least for Parkinson and Bornat) was to explain weak memory to ourselves. The observation that you don't seem to need isync/isb very often seems to have escaped most observers, but is perfectly clear in logical terms.

15.6 Loops

Loops, like conditionals, need special lacing. Not only do we have to deal with constraint coverage, we also have to cope with lo- and bo-parallel interference of assignments from successive executions of the loop body.

We illustrate the problems, and their solutions, by considering an MP receiver which loops, waiting to see $f = 1$, before it reads m. Unlabelled and unlaced we might write

$$\mathsf{do}\, r1 := f \,\mathsf{until}\, r1 = 1;\ r2 := m \tag{30}$$

[8] And we've heard a hardware architect say so

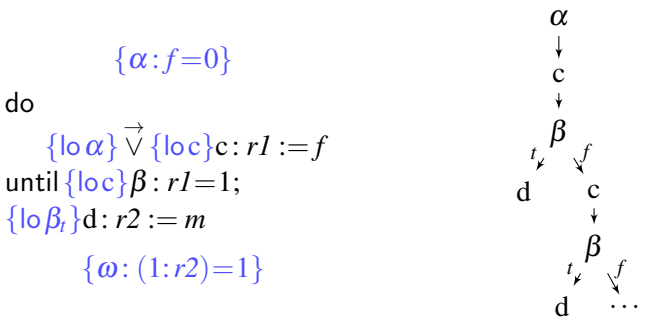

Figure 24: An MP receiver with a loop, and part of its so tree

Fig. 24 shows this receiver laced for weak memory, together with a prefix of its infinite so tree.

Consider first the tree. There is a path $\alpha \rightarrow c \rightarrow \beta$; from there via β_t we reach d, but via β_f we are back at c, and then on to β again. And so on, infinitely: each time round the loop there's a split into two paths. The infinite path which never takes the β_t branch contains both β and c infinitely often.

A non-disjunctive constraint $\alpha \rightarrow c$ in the program would have to stand for an infinite number of constraints in the tree, one to each instance of c, each carrying the same embroidery. That isn't impossible or even unreasonable, but it does produce a surprising result. A constraint which runs from α to the second instance of c, for example, is lo parallel with the first instance of c, which will interfere and may destabilise its embroidery. That's quite unlike the constraints which we've seen before: in the absence of loops, a constraint $a \rightarrow b$ is never lo parallel with either a or b.

What a component inside a loop requires is a stylised lacing which describes a constraint from the initial path to its first instance, and a constraint from that instance or something later in the loop to its next instance, and so on indefinitely. In the case of the looping receiver, the constraint on the first instance of c should come from α, and subsequently each instance should take its constraint from the previous one. The effect is an infinity of constraints, all but one identical, and none of them lo parallel with c. The overall constraint is the disjunction of the first constraint and one of the identical constraints.

All of that is indicated by the iterated constraint $\{\text{lo}\,\alpha\} \vec{\vee} \{\text{loc}\}$ in fig. 24; the first disjunct covers paths from the beginning of execution and the second covers paths round the loop. The rest of the lacing is straightforward: β always follows c; d always follows β_f; thread postcondition post, when we include one, always follows d. Note that each precondition has full coverage of the so-tree paths to its component.

Definition 9: Iterated constraint

A component x in a loop may have an iterated constraint $knot1 \vec{\vee} knot2$. Constraints in *knot1* must cover all paths from the initial assertion to x; constraints in *knot2* must cover all paths from x to x.

After all that preparation we discover that the proof in fig. 25 doesn't need to use an iterated constraint. The embroidery of the constraint $\alpha \rightarrow c$ doesn't mention *r1*, so is immune to the interference of c, so it can constrain all its instances. (We don't apologise for introducing the iterated constraint through this example: several later examples will need it.)

In the proof the precondition $f = 1 \Rightarrow m = 1$ of c is implied by initial assertion α, as before in fig. 9; it is stable against the rely, as before; it's stable against the interference of c itself (a new check); the constraint isn't lo-parallel with d, as before; and there is nothing to write in the guarantee, as before. The precondition of β is implied by the postcondition of c, is stable against the rely, and the constraint isn't lo-parallel with anything. The precondition of d is implied by the β_t postcondition $(r1 = 1 \Rightarrow m = 1) \wedge r1 = 1$. The postcondition of d

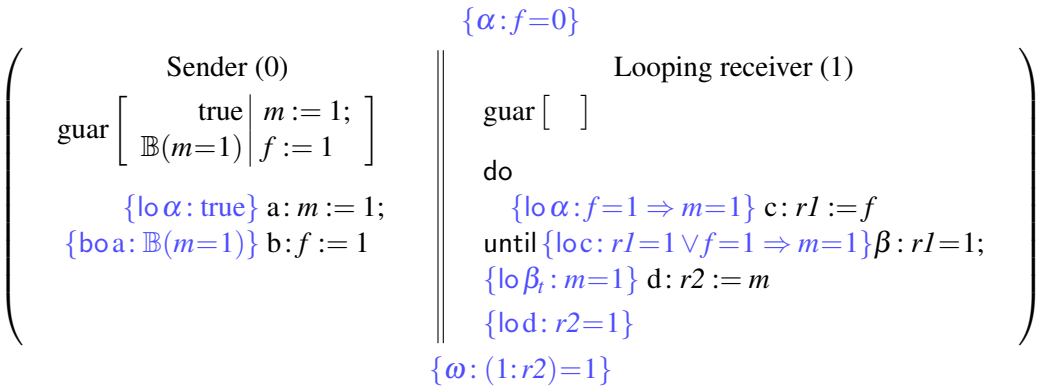
{α : f=0}
Sender (0)
guar [true | m := 1; ; B(m=1) | f := 1]
{lo α : true} a : m := 1;
{bo a : B(m=1)} b : f := 1
Looping receiver (1)
guar [ ]
do
{lo α : f=1 ⇒ m=1} c : r1 := f
until {loc : r1=1 ∨ f=1 ⇒ m=1} β : r1=1;
{lo β_t : m=1} d : r2 := m
{lo d : r2=1}
{ω : (1:r2)=1}

Figure 25: Proof of the looping receiver

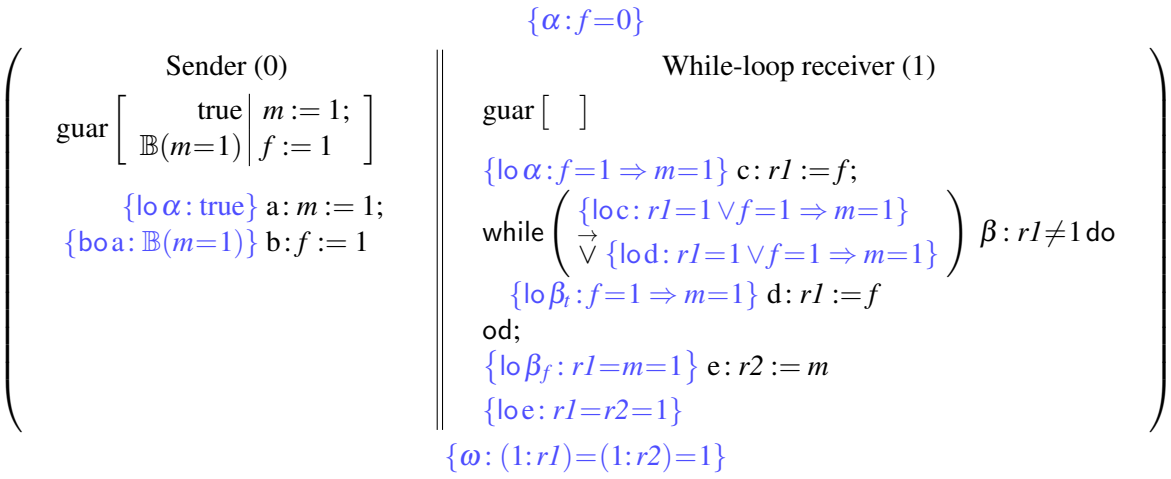
{α : f=0}
Sender (0)
guar [true | m := 1; ; B(m=1) | f := 1]
{lo α : true} a : m := 1;
{bo a : B(m=1)} b : f := 1
While-loop receiver (1)
guar [ ]
{lo α : f=1 ⇒ m=1} c : r1 := f;
while (→ {loc : r1=1 ∨ f=1 ⇒ m=1}) ∨ {lo d : r1=1 ∨ f=1 ⇒ m=1}) β : r1≠1 do
{lo β_t : f=1 ⇒ m=1} d : r1 := f
od;
{lo β_f : r1=m=1} e : r2 := m
{lo e : r1=r2=1}
{ω : (1:r1)=(1:r2)=1}

Figure 26: MP receiver with a while loop

implies $r2=1$, which is stable against anything, and the proof is complete.

In passing we note that this example shows why our logic can't, in general, deal with termination. We can't prove that the receiver will ever read 1 from f: the rely only says that the sender *might* sometimes write $f=1$.

We also note that, using the technique of section 15.5, this example doesn't need a barrier to implement the $\beta_t \rightarrow$d constraint. Loops just got cheaper (sometimes).

We used a do-until loop in our example because it avoids an invariant-intialising assignment and can illustrate both the possibility and the non-necessity of a loopback constraint. We can also handle a receiver with a while loop: see fig. 26. In this example we do need the loopback d$\rightarrow\beta$ in an iterated constraint, to ensure that components after the loop come after command d if it is elaborated. Without the loopback both the precondition of e and the thread postcondition are lo-unstable against the interference of d.

Part III

More precision and more orderings

16 Interference and stability

In our examples so far, all variable occurrences in interference preconditions have been $\mathbb{B}$-modal, and we have in effect used the SC stability rule 1 when dealing with external, lo-parallel or bo-parallel interference. But in general an interference precondition may contain amodal occurrences of variables, and then the SC rule won't work. It is time to confront the complexity of weak-memory stability checks and reveal the rules we actually use.

In this section we also discuss 'SC per location' aka 'global coherence', with which some, perhaps most, architectures constrain propagation reordering, and 'SC per register'. aka 'register renaming'. Each of these mechanisms affects internal and external parallelism and therefore our treatment of stability.

We begin, once again, by considering the stability of assertion P against interference $Q \mid x := E$, taking into account rule 6, repeated here for convenience.

> P is stable against $[\overline{ns}].Q \mid x := E$ if it is stable against $Q[\overline{ns}\backslash\overline{fs}] \mid x := E[\overline{ns}\backslash\overline{fs}]$ where $\overline{fs}$ are fresh names. Similarly for internal interference $[\overline{ns}].Q \mid r := E$ and $[\overline{ns}].Q \mid r := x$.

16.1 Internal and external parallelism

In SC execution the interfered-with P-asserting thread and the interfering Q-asserting thread see the same state of the store, so rule 1 considers the case when both P and Q hold. For *internal* (lo-parallel) stability, coming from an assigment in the same thread, both assertion and assignment operate within the same local 'store' of buffers and caches so we can, in effect, use the SC rule – that is, rule 4 of section 11.1, repeated here for convenience.

> Constraint embroidery P is LO stable against lo-parallel assignment A with precondition Q if
> $$\mathrm{sp}(P \wedge Q,\, A) \Rightarrow P$$

Internal interference, just like external interference, is quotiented to hide registers (see section 17.3 for further discussion).

When interference comes from another thread we can't suppose that in general Q must hold in the P-asserting thread at the instant when $x := E$ interferes. Operationally this is because $x{=}E$ might arrive before all or some of the writes which underpin Q. In previous examples we have used $\mathbb{B}$-modal Q to eliminate propagation delay. Now we consider the writes which underpin amodal occurrences of variables in Q. Our treatment of stability must ensure that P is stable in all the possible states corresponding to early or late propagation of those writes. We do that by *hatting* amodal occurrences of variables to decouple them from the values of the same variables if they occur in P:

Rule 7: EXT stability

> Constraint embroidery P is EXT stable against $Q \mid x := E$ from another thread if
> $$\mathrm{sp}(P \wedge \widehat{Q},\, x := E) \Rightarrow P$$

$\widehat{Q}$ is a copy of Q in which every free occurrence of a variable v is renamed to $\widehat{v}$, *except* those inside a $\mathbb{B}$ or $\mathbb{U}$ modality.

Definition 10: Hatting

$$
\begin{array}{rcll}
\widehat{var} & = & \widehat{var} & \text{whether free or bound} \\
\widehat{reg} & = & reg & \\
\widehat{const} & = & const & \\
\widehat{P \text{ binop } Q} & = & \widehat{P} \text{ binop } \widehat{Q} & \\
\widehat{\text{unop } P} & = & \text{unop } \widehat{P} & \\
\widehat{\text{binder } v(P)} & = & \text{binder } v(\widehat{P}) & \\
\widehat{\mathbb{B}(P)} & = & \mathbb{B}(P) \wedge \widehat{P} &
\end{array}
$$

The intuition is that v means 'v in the asserting thread, when the interference arrives' and $\widehat{v}$ means 'v in the interfering thread, when the interference was elaborated'. Note, then, that hatting operates even within binders: $\widehat{\forall x(P(x))}$ is a remark about variables that held when the interference was constructed and is adequately described by $\forall x(P(\widehat{x}))$ if we recognise that the binding quotients the identity of the variable, not its temporal or thread-specific occurrences. Expansion of the $\mathbb{B}$ modality allows the enclosed assertion to apply both when the interference arrives and when it was elaborated.

Observe that hatting separates the $\mathbb{B}$-modal parts of Q, whose underpinning writes must be propagated before $x{=}E$, from the rest of Q, whose writes may or may not be propagated before $x{=}E$. Retaining the amodal writes in hatted form enables us to continue to use E, which may refer via registers to those hatted occurrences, to describe the value assigned by the interference.

Hatting decouples amodal variable occurrences in Q from similar occurrences in P. To impose some correspondence between hatted and unhatted versions – say $y{=}\widehat{y}$ – would merely strengthen the left-hand side of the implication in rule 7. So if P is stable when all amodal writes arrive late, then it is stable when some arrive on time.

Previous examples had no amodal variable occurrences in interference preconditions, so the interference precondition Q was preserved by hatting, and the extra hatted formulas derived from $\mathbb{B}$ modalities were irrelevant in our examples, in which the assigned expression E didn't depend on the values of variables.

16.2 In-flight interference

The EXT stability rule 7 makes it clear that the important part of an interference precondition is its unhatted – modal – part. We might suppose, then, that the in-flight stability of $P \mid y := F$ against bo-parallel $Q \mid x := E$ could be checked by

$$\text{sp}(\widehat{P} \wedge Q,\ x := E) \Rightarrow \widehat{P} \tag{31}$$

But that would provide trivial stability in the case where the amodal parts of Q contradict the modal parts of P. That can happen if P, modally, and Q, amodally, mention x, and/or if lo-intervening assignments change other variables mentioned modally in P and amodally Q. So we need to decouple the modal parts of P from the amodal parts of Q. We can't simply hat P, because that would couple the amodal parts of P and Q. Instead we hat P with a different accent: $\widehat{\widehat{P}}$ is just like $\widehat{P}$ in its operation.

Rule 8: In-flight (bo) stability

Interference $P \mid y := F$ is BO stable against bo-parallel $Q \mid x := E$ if

$$\text{sp}(\widehat{P} \wedge \widehat{\widehat{Q}},\ x := E) \Rightarrow \widehat{P}$$

In this rule there are unhatted variables, whose values are propagated before $y{=}F$ and $x{=}E$; hatted variables,

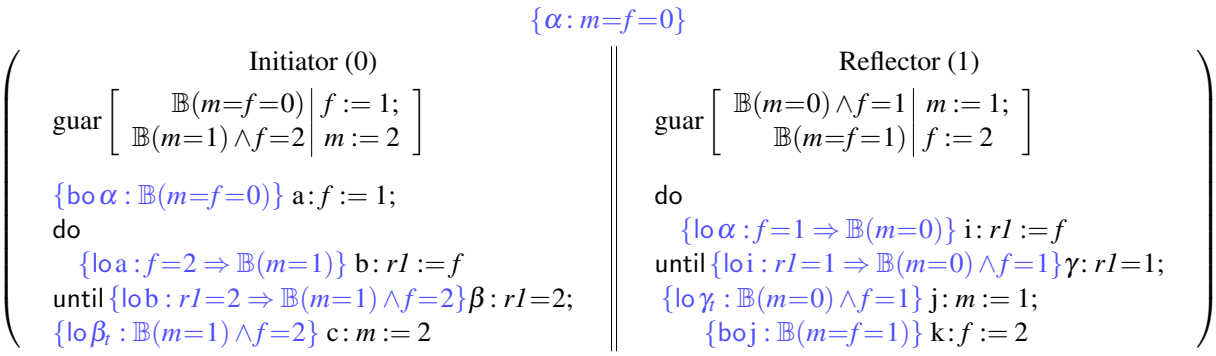
{α : m=f=0}
Initiator (0)
guar [B(m=f=0) | f := 1; ; B(m=1)∧f=2 | m := 2]
{bo α : B(m=f=0)} a: f := 1;
do
{lo a : f=2 ⇒ B(m=1)} b: rl := f
until {lo b : rl=2 ⇒ B(m=1)∧f=2} β : rl=2;
{lo β_t : B(m=1)∧f=2} c: m := 2
Reflector (1)
guar [B(m=0)∧f=1 | m := 1; ; B(m=f=1) | f := 2]
do
{lo α : f=1 ⇒ B(m=0)} i: rl := f
until {loi : rl=1 ⇒ B(m=0)∧f=1} γ: rl=1;
{lo γ_t : B(m=0)∧f=1} j: m := 1;
{boj : B(m=f=1)} k: f := 2

Figure 27: Spurious bo parallelism

whose values applied when P was elaborated; and double-hatted variables, whose values applied when Q was elaborated.

This rule applies to the interference of assignments within a thread which are bo parallel: that is, assignments to different variables which are so but not bo ordered. It also applies to interferences in the rely that come from different threads and assign to different variables: in that case we can think of x as 'x in the interfered-with thread', $\widehat{x}$ as 'x in one interfering thread' and $\widehat{\widehat{x}}$ as 'x in the other interfering thread'.

BO parallelism asymmetrical within a thread, according to definition 6, repeated here for convenience.

> If variable assignments to distinct variables are so- but not bo-ordered then the so-later is *bo parallel* with the so-earlier and will interfere with its interference precondition.

The elaboration precondition of the later assignment will have, baked in, the effect of the first assignment.

But why should we take any account at all of the $\mathbb{B}$-modal parts of Q in the BO stability check? The main reason is that the test for bo parallelism in definition 6 is too weak: it doesn't take account of what is called 'detour' in (Alglave et al., 2014): causal chains between two assignments which may lead through more than one thread back to the first and include a bo in one of the threads it passes through. In fig. 27, for example, there is a chain a→i→γ→j→k→b→β→c. Execution must traverse that path: after a we have $f=1$ and the value of f won't change without the action of the reflector thread. The j→k constraint is bo, and in the proof, even though a and c are bo parallel according to definition 6, they have contradictory $\mathbb{B}$-modal interference preconditions and so c won't destabilise the interference of a.

16.3 A summary of stability checks

Internal stability (within a thread): LO rule 4.

External stability (inter-thread): EXT rule 7.

Stability of interference preconditions (within a thread; between threads in the rely): BO rule 8.

There remain UEXT and UO stability, which deal with $\mathbb{U}$-modal and $\mathbb{S}$ofar assertions and are discussed in section 19.

17 Implicit constraints

In our characterisation of weak memory, reordering of command executions within a thread applies to a path in the so tree and is otherwise limited only by programmer-supplied constraints. In practice hardware

architectures supply constraints of their own, but because our logic uses embroidery on explicit constraints, we can't easily take advantage of them. That isn't a problem: superfluous constraints compile to no code, and the absence of a constraint allows but does not enforce parallelism.

But there are two constraints which seem to be almost universally provided, and which therefore might be exploited: ordering of writes to a single location, and register renaming.

17.1 SCloc (global coherence)

In many weak-memory architectures, including the major CPU architectures of x86, Power and ARM, there is a global constraint on the propagation of writes. If one thread 'witnesses' two writes to a location in a particular order – say $x{=}1$ before $x{=}2$ – then no other thread will 'witness' those writes in the opposite order. 'Witnessing' covers both reading and writing, and the constraint doesn't require that other threads will witness those writes at all: only that if they are witnessed then they are witnessed in the same order. In detail it's more subtle than that, but in a logic in which nothing is in order unless it is lo (or bo or uo or go) constrained, it can be understood.

The constraint is called *global coherence* in the Power and ARM architectures, and *SC per location* in (Alglave et al., 2014); for brevity we call it SCloc. Under SCloc each thread witnesses, for each location, a subset of the partial order of all writes of all threads to that location.

It turns out (see section 26) that termination reasoning in our logic depends on SCloc, and we believe that our use of the $\mathbb{B}$ and $\mathbb{U}$ modalities would be unsound without it. But all that is implicit: the visible effects of SCloc on our logic are only two:

1. Successive writes to the same location within a single thread will not be propagated out of lo order. So the definition of bo parallel in definition 6 applies only to assignments to different variables.

2. We can reason about the overall coherence order of individual variables (see section 25).

SCloc is subtle. It's by no means true, for example, that on every architecture a write to a variable is always lo-ordered with a so-later read from the same variable. So we assume no implicit ordering, relying instead on explicit constraints.

17.1.1 (P) RDW and CoRR litmus tests

On Power and ARM the SCloc constraint demands that when two so-ordered reads from the same location take their values from two different external writes, there is an implicit elaboration order between them – i.e. an implicit lo constraint. This principle is exploited in the CoRR0/1/2, RDW and RSW litmus tests of (Pow, 2015; Maranget et al., 2012). But it only applies when reading from *external* writes, those generated in other threads. Maranget et al. (2012) give a litmus test which shows that, on ARM at least, a read from an external write is not necessarily lo-ordered with an so-later internal write to the same variable (and we know of other more subtle examples which illustrate similar pecularities).

In a logic based on state assertions it would be difficult to make a distinction between internal and external writes, and we haven't tried to make one. RDW and RSW, which appeal to an implicit dependency, are therefore beyond us. But by including explicit lo constraints in the observer threads of CoRR0/1/2, we can deal with the principle behind those tests. Our treatment is illustrated in several examples in section 25 and subsequently.

$$\text{guar}\left[\begin{array}{l|l} \text{true} & y := 1; \\ \text{true} & x := 2 \end{array}\right]$$

$$\begin{array}{r} \text{a}: r1 := 1; \\ \{\mathsf{lo}\,\text{a}: r1{=}1\}\ \text{b}: y := r1; \\ \text{c}: r1 := 2; \\ \{\mathsf{lo}\,\text{c}: r1{=}2\}\ \text{d}: x := r1 \\ \{\mathsf{lo}\,\text{d}: x{=}2; \mathsf{lo}\,\text{b}: y{=}1\} \end{array}$$

Figure 28: Valid proof with SCreg, invalid without

17.1.2 A retreat

In earlier versions of this work we made an argument from eventual consistency and SCloc that an assertion could not experience interference from an external assignment that originated from a contradictory stable state. That enabled us to strengthen the EXT stability check and permitted thread-local proofs of the LB litmus tests (fig. 4) and the token ring (section 30). We cannot maintain that argument; amongst other things, it made it impossible to have a simple treatment of BO stability. We have new proofs of those tests, but we have to use more powerful machinery: see section 23 and section 30.

17.2 Register renaming (SCreg)

Our logic treats registers as local variables, invisible to other threads, and uses the LO stability rule 4 to deal with lo parallelism of register assignments and assertions within a thread. Most modern architectures take a different approach, in which registers are more like single-assignment variables, and it makes a difference. The mechanism enables thread-local speculation of register assignments, and so it seems to be used in every architecture.

Consider, for example, the single-thread program in fig. 28. According to our treatment so far, commands a and b are lo parallel with the c→d constraint (so-before but not lo-before), and c and d are lo-parallel with a→b (so-after but not lo-after). So by the rules of our logic a interferes with the $r1{=}2$ embroidery on the c→d constraint, and c interferes with the $r1{=}1$ embroidery of the a→b constraint. Each embroidery is obviously unstable in those circumstances. But in a machine which implements register assignment by taking a register from a pool and assigning it a value, and using that new register for so-later commands until the next assignment to the 'same' register, the *r1* in commands a and b is not actually the same register as the *r1* in commands c and d, so the lo parallelism doesn't cause instability in either case.

Under *register renaming* (SCreg) each reference to a register takes its value from the so-latest assignment to that register and there can be no interference from other assignments using the same register name. That makes register interference very much ruled by so. We can make our logic do the same thing if we said that lo parallelism of a register assignment and a constraint requires the assignment to be so-between the constraint's source and target. If we do that, our Arsenic proofchecker accepts the proof in fig. 28.

In this paper we don't exploit SCreg, because our examples don't need it. Even if we did exploit it, we would still require explicit lo constraints between a register assignment and subsequent uses of the register, because constraints carry embroidery, and embroidery is essential to proof. But since those constraints are free on SCreg machinery, we don't need to fear their cost.

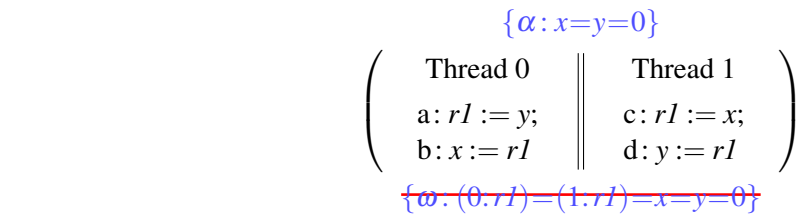

Figure 29: 42 from thin air?

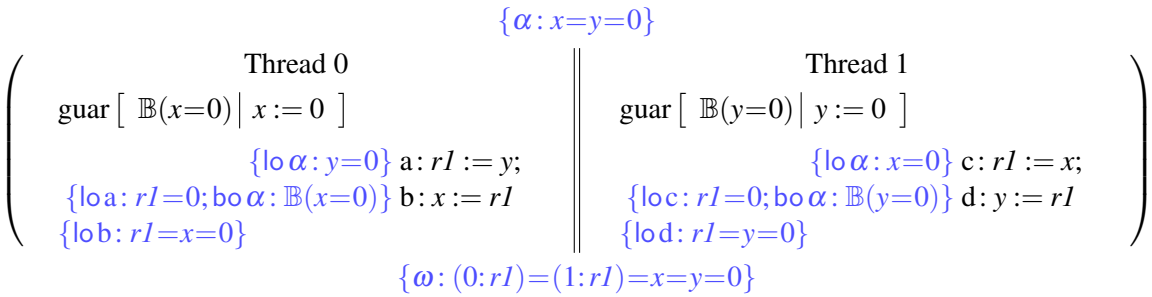
{α : x=y=0}
Thread 0
guar [𝔹(x=0) | x := 0]
{lo α : y=0} a : r1 := y;
{lo a : r1=0; bo α : 𝔹(x=0)} b : x := r1
{lo b : r1=x=0}
Thread 1
guar [𝔹(y=0) | y := 0]
{lo α : x=0} c : r1 := x;
{lo c : r1=0; bo α : 𝔹(y=0)} d : y := r1
{lo d : r1=y=0}
{ω : (0:r1)=(1:r1)=x=y=0}

Figure 30: Nothing from thin air with SCreg

17.3 Quotienting internal interference

In checking internal interference – LO, BO and UO – we have to be careful with registers. Because of SCreg the same register name in different assertions may not refer to the same register. When dealing with inter-thread interference (EXT and UEXT) we quotient register names because threads don't share registers. In the absence of an analysis which could show whether two occurrences of the same register name refer to the same register, we quotient internal interference in just the same way that we do external interference. This causes considerable incompleteness (sob).

17.4 An example that appeals to SCreg dependencies

In the unconstrained program of fig. 29, we imagine that weak-memory hardware can choose the elaboration order of each thread. A machine might choose 42 as the initial value of $r1$ in thread 0 and, elaborating and propagating b first, write x=42. Then thread 2 might elaborate in so order, reading x=42 and writing y=42. Finally command a of thread 0 could read y=42. It hasn't followed so order in thread 1, but it has used a reordering of the so tree, so there is an execution which terminates with all of its registers and all of its variables containing 42. Given an unconstrained program it's quite possible and perfectly understandable: in weak memory we can't be sure to finish with x=y=0.

But suppose we execute this machine on SCreg hardware: then there would be implicit lo constraints a→b and c→d. Fig. 30 makes those constraints explicit, and shows that neither thread generates any effective interference, each writing 0 to a variable that is already 0.[9] There's no possibility, under SCreg, of a cycle in which both threads read and write a value plucked out of 'thin air' (or rather, taken from an uninitialised

[9] This proof is a good example of circular rely/guarantee reasoning: thread 0 only writes 0 because thread 1 only writes 0 because thread 0 only writes 0 because ...

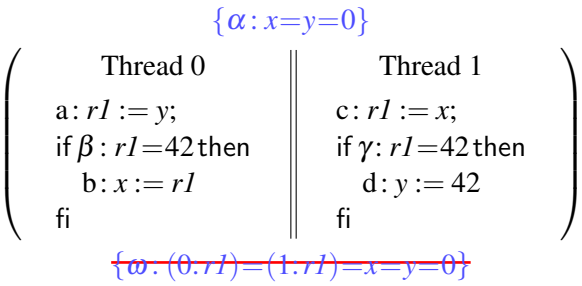

Figure 31: 42 from slightly thicker air?

register).

The constrained program works whether or not there is SCreg, but on an SCreg machine the lo constraints don't need implementation. On such machines there is no 'thin air' execution of the apparently unconstrained fig. 29.

18 Speculated propagation and the go ordering

In fig. 31, adapted from the C11 standard (C, 2011), there is a command which writes 42 so – unlike fig. 29 – the hardware doesn't have to make it up. The standard doesn't prohibit an execution which finishes with $x{=}y{=}42$, but its authors express a hope that compiler writers will conspire to make it impossible.[10] It seems a bit of a nightmare.

To see how the nightmare could be real we trace a possible execution of the unconstrained program. Command d in thread 1 is elaborated first and $y{=}42$ propagated; then thread 0 executes in so order, reading $y{=}42$, finding $r1{=}42$ and writing $x{=}42$; then c of thread 1 reads 42 and control expression γ confirms that d indeed ought to be elaborated. That execution obeys the constraint on execution of threads that we have assumed throughout: each thread's execution follows a reordering of an so-tree path. So perhaps the nightmare is real. If it is real it's important: Batty et al. (2011) show that 'satisfaction cycles', in which a command's execution is the cause of itself, prevent effective library abstraction in C11. In our nightmare the elaboration of $y := 42$ is eventually a cause, via thread 0, command c and control expression γ, of its own presence in the so path used by thread 1.

But like most nightmares there's something odd about it. Suppose that d is speculatively elaborated and $y{=}42$ propagated from thread 1, but that thread 0 is quick off the mark, reading $y{=}0$ from the initial state before $y{=}42$ reaches it. Then thread 0 won't write $x{=}42$, and therefore c must read $x{=}0$ from the initial state, γ will find $r1{\neq}42$ and execution must follow the γ_f path, on which command d doesn't appear. So in those circumstances thread 1 has *not* followed a reordering of an so-tree path. It would either have to withdraw the propagation of $y{=}42$ and retreat to an execution in which it wrote nothing, or negotiate with thread 0 to re-execute reading 42 and then recapitulate its own execution of c and γ.

This example shows that it can require inter-thread cooperation to allow withdrawal of speculated propagation. Such cooperation is difficult if one or more threads has read from and acted upon the propagation, as thread 0 did in our imagined execution. So far as we aware, and perhaps as a consequence of the difficulty, no architecture attempts speculated propagation: they speculate only elaboration, which is internal to

[10] The appeal is to compiler writers and not to hardware designers, and that's important. Our treatment is of a kind of high-level assembler. The problems of defining a programming language are different: see section 33.

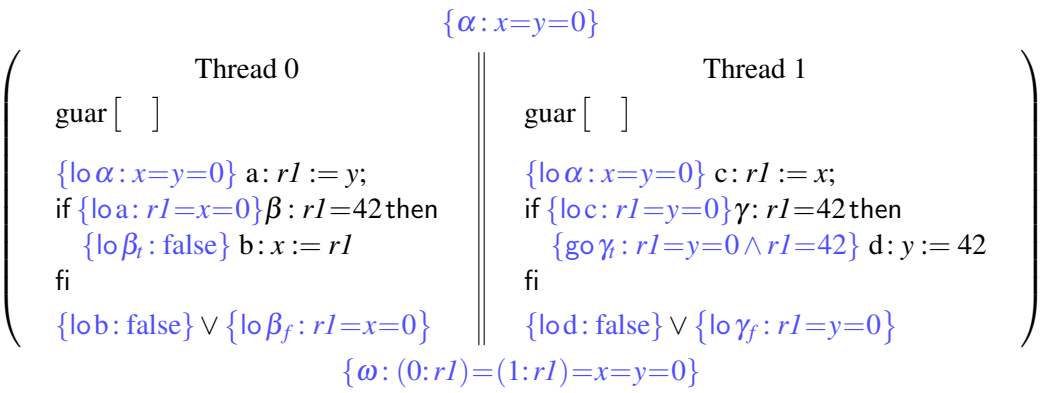
{α : x=y=0}
Thread 0
guar [ ]
{lo α : x=y=0} a : r1 := y;
if {lo a : r1=x=0} β : r1=42 then
{lo β_t : false} b : x := r1
fi
{lo b : false} ∨ {lo β_f : r1=x=0}
Thread 1
guar [ ]
{lo α : x=y=0} c : r1 := x;
if {lo c : r1=y=0} γ : r1=42 then
{go γ_t : r1=y=0 ∧ r1=42} d : y := 42
fi
{lo d : false} ∨ {lo γ_f : r1=y=0}
{ω : (0:r1)=(1:r1)=x=y=0}

Figure 32: No thin air with SCreg dependencies and LocalSpec

a thread and which, because of register renaming, is relatively easily abandoned.

To deal with this problem, we add LocalSpec to the architecture parameters SCloc and SCreg. Under LocalSpec all speculation is local – i.e. no speculated propagation of writes. A machine without LocalSpec can produce the execution we described. One with SCreg and LocalSpec cannot, as the proof in fig. 32 shows. We have made explicit the SCreg-implicit lo constraints a$\rightarrow\beta_t\rightarrow$b and c$\rightarrow\gamma_t$. We have also made explicit the LocalSpec-implicit go constraint $\gamma_t\rightarrow$d.

The target of a go contraint must be a variable assignment, and the *propagation* of the write generated by the target is constrained to occur after the *elaboration* of the source. So go doesn't fit in the lo/bo/uo hierarchy, and in particular go is not included in lo. In terms of proof a go constraint's embroidery contributes primarily to the interference rather than the elaboration precondition of its target. The go constraint in fig. 32 provides an interference precondition of d which is false – i.e. it won't be propagated, ever.

Execution follows a reordering of an so-tree path. A path cannot contain an assignment that cannot be propagated, so the elaboration precondition of d must be false as well. To provide this, the rules for forming a precondition are as follows:

Definition 11: Precondition from knot

- The overall precondition is the conjunction of a knot's embroidery or, if the knot is a disjunction of knots, the disjunction of their overall preconditions.
- The elaboration precondition is the overall precondition, excluding the embroidery of go constraints, conjoined with the satisfaction (sat) of the overall precondition.
- The interference precondition is either the overall precondition or, if the knot contains a square-bracketed *intfpre* assertion (see table 2), the *intfpre* assertion, which must be implied by the overall precondition.

If there are no go constraints in a knot then the elaboration precondition is just the overall precondition. But in the case of command d in fig. 32 we must exclude the embroidery of the go constraint and conjoin the satisfaction of the overall precondition sat($r1=y=0\wedge r1=42$), which is false. An assignment with a false precondition needs no guarantee entry and causes no lo-parallel or bo-parallel instability; we have a proof.

Without the go constraint the proof in fig. 32 fails spectacularly. The preconditions of command c and control expression γ are lo-parallel with command d, and therefore their embroidery is unstable.[11] The

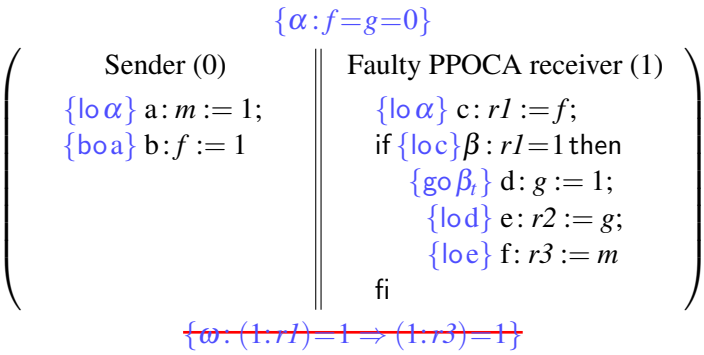

Figure 33: PPOCA: ineffective message-passing

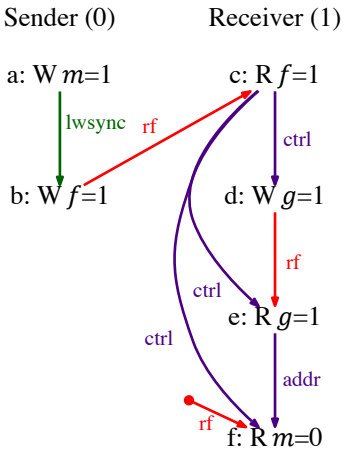

Figure 34: PPOCA litmus test diagram

guarantee should include true | $y := 42$; and the embroidery of the d→thread-postcondition constraint can't be false.

We can't prove that the nightmare execution can't happen on any weak-memory architecture; without LocalSpec our proof fails to show that it can't happen; and indeed we believe that it could happen in that case. But we can prove that it can't happen with commonly-applied architectural constraints – with LocalSpec and with SCreg.

18.1 (P) PPOCA: a scary conditional

Fig. 33 shows a version of MP with a receiver which conditionally writes and reads g before reading m: note that the constraint β_t→d is go rather than lo. It doesn't work, just because go doesn't imply lo. Fig. 34 shows an observed weak-memory execution with a chain of various implicit and explicit Power-specific orderings from command c to command f which don't compose to produce the equivalent of an lo constraint. Experimentally the receiver can read 1 from f and then 0 from m.

The issue is the difference between elaboration and propagation. In the unproof of fig. 35 we have repro-

[11] They are lo-parallel with the go constraint in place, but the false elaboration precondition means no instability.

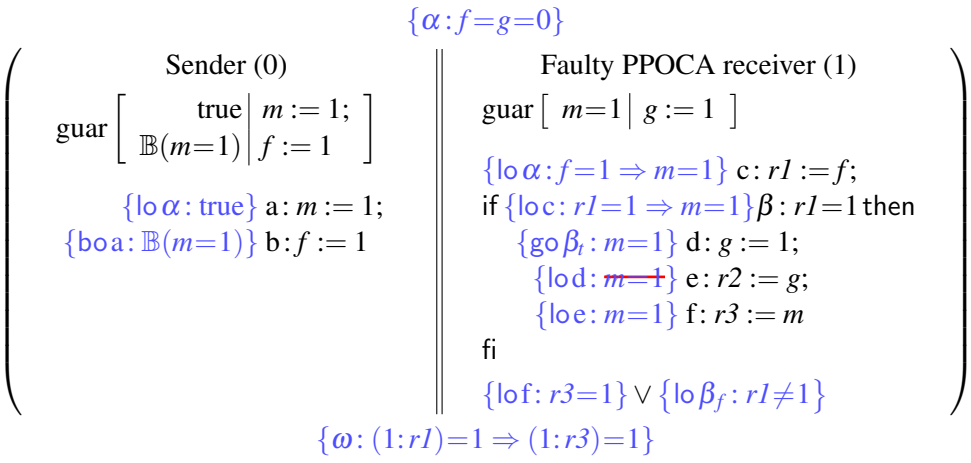
{α : f=g=0}
Sender (0)
guar [true | m := 1; B(m=1) | f := 1]
{lo α : true} a : m := 1;
{boa : B(m=1)} b : f := 1
Faulty PPOCA receiver (1)
guar [m=1 | g := 1]
{lo α : f=1 ⇒ m=1} c : r1 := f;
if {loc : r1=1 ⇒ m=1} β : r1=1 then
{go β_t : m=1} d : g := 1;
{lod : m=1} e : r2 := g;
{loe : m=1} f : r3 := m
fi
{lof : r3=1} ∨ {lo β_f : r1≠1}
{ω : (1:r1)=1 ⇒ (1:r3)=1}

Figure 35: PPOCA unproof

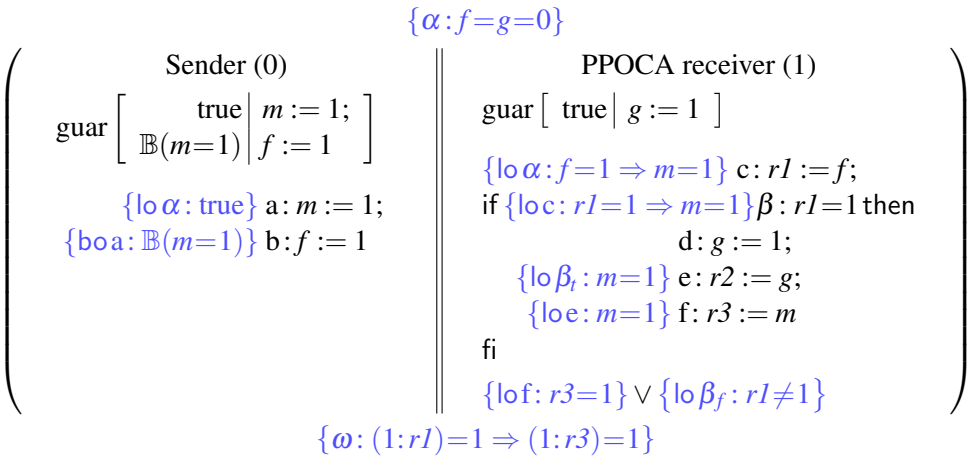
{α : f=g=0}
Sender (0)
guar [true | m := 1; B(m=1) | f := 1]
{lo α : true} a : m := 1;
{boa : B(m=1)} b : f := 1
PPOCA receiver (1)
guar [true | g := 1]
{lo α : f=1 ⇒ m=1} c : r1 := f;
if {loc : r1=1 ⇒ m=1} β : r1=1 then
d : g := 1;
{lo β_t : m=1} e : r2 := g;
{loe : m=1} f : r3 := m
fi
{lof : r3=1} ∨ {lo β_f : r1≠1}
{ω : (1:r1)=1 ⇒ (1:r3)=1}

Figure 36: PPOCA effectively laced

duced the orderings of the original: bo a→b mimics Power lwsync; go β_t→d mimics the Power control dependency; lo d→e is implicit on Power; lo e→f mimics the address dependency of the original.[12]

In fig. 35 the elaboration precondition of e, according to definition 11, is sat(m=1), which doesn't imply m=1: so the postcondition doesn't imply m=1; so we can't embroider m=1 on the lo constraint d→e; so it's an unproof. But note that the interference precondition of d is m=1, as recorded in the guarantee.

Some litmus tests are subtle, but PPOCA frightens the horses, because it seems at first sight that conditionals don't do what they ought to. Once we realise that go doesn't imply lo the solution is obvious: use an lo constraint either to e, as in the proof of a corrected PPOCA in fig. 36, or direct to f. We don't need to lace d because we don't care when it's elaborated. The horses may relax; the ordering they seek can easily be provided.

[12] We haven't mimicked Power's control dependencies to register assignments e and f, because go constraints can't target register assignments. But then control dependencies to register assignments don't have any force in the Power memory model either.

Part IV
Temporal reasoning and coherence

Our treatment thus far has implicitly used temporal reasoning. Our $\mathbb{B}$ modality employs a specialised version of the temporal 'since' modality: that fact is implicit in the treatment of substitution and $\mathbb{B}$ (definition 3). But we need other modalities to be able to deal with many of the features of weak memory: x86's MFENCE, Power's sync and ARM's dsb need a modality which applies across threads; 'coherence' arguments need modalities which describe the global temporal order of writes; and several litmus tests are about events that occur during the period when a particular property holds.

All of our temporal reasoning is based on 'since': an assertion has held – locally or globally – since a barrier event or since execution started; or an assertion held at least once since execution started.

Temporal modalities are very useful, but they introduce a difficulty: assertions which employ them can't simply be treated as descriptions of the current state of a thread. This undermines the justification of the $\mathbb{B}$ modality given in aside 1, and affects the treatment of termination. We deal with those difficulties in part V.

As usual we introduce our treatment with examples.

19 A proof of SB using $\mathbb{U}$ and since

In introducing weak memory in section 8, our third litmus test in fig. 6 was SB, abstracted from Dekker's algorithm for mutual exclusion. In weak memory, without constraints, both threads can read 0 as pictured in fig. 7. We can't make the program work by using lo or bo constraints: lo has no effect on propagation, so both reads could happen before the writes propagate; bo would do no more because the reads don't cause propagation. And in any case experiment shows that it takes a uo-generating barrier in each thread to make the program work.

Fig. 37 shows a laced version of SB, with uo constraints a→b and c→d. Uo provides lo and bo ordering, but does more than bo in that the writes which underpin the source postcondition are propagated to all threads before the target elaborates. So every thread's cache contains $x{=}1$ before thread 0 reads y from its own cache, and every thread's cache contains $y{=}1$ before thread 0 reads x. From the point of view of each thread we've established something like SC interleaving: the effect of a takes place everywhere before b is elaborated, and likewise for c before d.

Along with the uo ordering goes a $\mathbb{U}$ modality, and like the $\mathbb{B}$ modality on a bo constraint, we wrap an assertion implied by the source postcondition in $\mathbb{U}$ to get the constraint embroidery – see rule 2.

$\mathbb{U}(P)$ asserts that P holds in all threads since a barrier event; when we introduce it via a uo constraint the barrier event happens after the elaboration of the constraint's source and before the elaboration of its target. In the proof of SB in fig. 38, we claim that $x{=}1$ holds in all threads when command b is elaborated. If b then reads 0 from y, there was therefore an instant in thread 0 in which simultaneously $x{=}1$ held in every thread and $y{=}0$ held in thread 0. $\mathbb{U}(x{=}1)\,\mathsf{since}\,y{=}0$ is a temporal modality which claims that the coincidence occurred and that $\mathbb{U}(x{=}1)$ has held continuously ever since. Similarly in thread 1 we claim

$$\{\alpha : x{=}0 \wedge y{=}0\}$$

Thread 0	Thread 1
$\{\mathsf{lo}\,\alpha\}$ a: $x := 1$;	$\{\mathsf{lo}\,\alpha\}$ c: $y := 1$;
$\{\mathsf{uo}\,\mathrm{a}\}$ b: $r1 := y$	$\{\mathsf{uo}\,\mathrm{c}\}$ d: $r1 := x$

$$\{\omega : \neg((0{:}r1){=}0 \wedge (1{:}r1){=}0)\}$$

Figure 37: SB laced for weak memory with uo constraints

$$\{\alpha : x{=}0 \wedge y{=}0\}$$

$$\left(\begin{array}{c|c} \text{Thread 0} & \text{Thread 1} \\ \text{guar}\left[\ \text{true}\ \middle|\ x := 1\ \right] & \text{guar}\left[\ \text{true}\ \middle|\ y := 1\ \right] \\ \{\mathsf{lo}\,\alpha : \text{true}\}\ \mathrm{a} : x := 1; & \{\mathsf{lo}\,\alpha : \text{true}\}\ \mathrm{c} : y := 1; \\ \{\mathsf{uo}\,\mathrm{a} : \mathbb{U}(x{=}1)\}\ \mathrm{b} : r1 := y & \{\mathsf{uo}\,\mathrm{c} : \mathbb{U}(y{=}1)\}\ \mathrm{d} : r1 := x \\ \{\mathsf{lo}\,\mathrm{b} : r1{=}0 \Rightarrow (\mathbb{U}(x{=}1)\,\mathsf{since}\,y{=}0)\} & \{\mathsf{lo}\,\mathrm{d} : r1{=}0 \Rightarrow (\mathbb{U}(y{=}1)\,\mathsf{since}\,x{=}0)\} \end{array}\right)$$

$$\{\omega : \neg((0{:}r1){=}0 \wedge (1{:}r1){=}0)\}$$

Figure 38: A weak-memory proof of SB

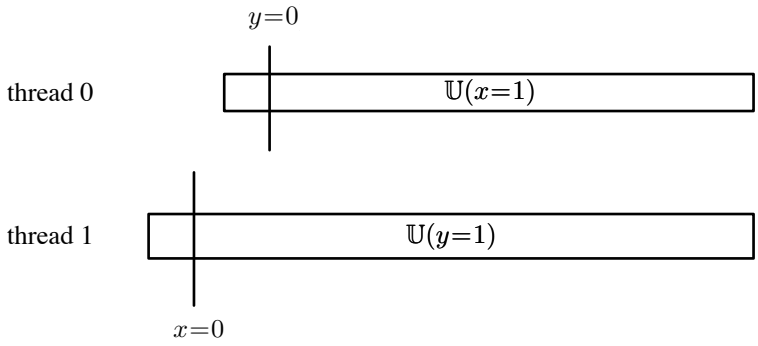

Figure 39: An impossible SB thread history

$\mathbb{U}(y{=}1)\,\mathsf{since}\,x{=}0$. The thread-history diagram in fig. 39 illustrates why those claims can't both be true: if the $\mathbb{U}(x{=}1)$ epoch starts late enough for thread 1 to see $x{=}0$ when $y{=}1$ already holds everywhere, then thread 0 can't later see $y{=}0$ – and a similar diagram would show that if thread 0 sees $y{=}0$ then thread 1 can't see $x{=}0$. Mutual exclusion, as required, but so far only an informal proof.

Since the two threads of SB are similar, swapping the rôles of x and y, we check only thread 0 of the proof in fig. 38. The precondition true of command a is stable against any interference, and a gives us the guarantee of thread 0. The command's postcondition implies $x{=}1$, which is stable because thread 1 writes only to y; we decorate that postcondition with $\mathbb{U}$ in a uo-ordering to b.

To deal with the pre and post conditions of b we have to consider the formal properties of $\mathbb{U}$. It shares all the properties in definition 4, reading $\mathbb{U}$ for $\mathbb{B}$, plus the following:

Definition 12: Additional properties of the $\mathbb{U}$ modality

$$\begin{array}{rcl} \mathbb{U}(P) & \Rightarrow & \mathbb{B}(P) \\ \mathbb{U}(\mathbb{B}(P)) & = & \mathbb{U}(P) \\ \mathbb{B}(\mathbb{U}(P)) & = & \mathbb{U}(P) \end{array}$$

Substitution and hatting with $\mathbb{U}$ is just as with $\mathbb{B}$.

Definition 13: Substitution and hatting of the $\mathbb{U}$ modality

$$\begin{array}{rcl} \mathbb{U}(P)[x\backslash\overleftarrow{x}] & = & \overleftarrow{\mathbb{U}(P)} \wedge P[x\backslash\overleftarrow{x}] \\ \overleftarrow{\mathbb{U}(P)} \wedge P & \Rightarrow & \mathbb{U}(P) \\ \mathbb{U}(P)[r\backslash\overleftarrow{r}] & = & \mathbb{U}(P[r\backslash\overleftarrow{r}]) \\ \widehat{\mathbb{U}(P)} & = & \mathbb{U}(P) \wedge \widehat{P} \end{array}$$

So the precondition $\mathbb{U}(x{=}1)$ of b is stable against the interference of thread 1 because thread 1 doesn't interfere with x. The instantaneous postcondition of b is

$$\mathbb{U}(x{=}1) \wedge r1{=}y \tag{32}$$

which we may weaken to

$$r1{=}0 \Rightarrow (\mathbb{U}(x{=}1) \wedge y{=}0) \tag{33}$$

This says that when $r1{=}0$ thread 0 has witnessed the case we are interested in, in which it wrote 1 and read 0. But this assertion is unstable because thread 1 can write $y{=}1$. We further weaken to

$$r1{=}0 \Rightarrow (\mathbb{U}(x{=}1)\,\mathsf{since}\,y{=}0) \tag{34}$$

which is stable against thread 1's interference despite the fact that it mentions $y{=}0$. That's the magic of 'since'. To understand it we need to look at the formal properties of 'since'.

$P\,\mathsf{since}\,Q$ holds if $P \wedge Q$ holds, or if $P \wedge Q$ held in the past and P has held ever since. It's established by a temporal coincidence; it implies P; and it implies that Q happened, which we capture with the $\mathbb{O}\mathsf{uat}$ – 'once upon a time' – modality, to be discussed in section 24.

Definition 14: Properties of since

$$\begin{array}{rcl} P \wedge Q & \Rightarrow & P\,\mathsf{since}\,Q \\ P\,\mathsf{since}\,Q & \Rightarrow & P \wedge \mathbb{O}\mathsf{uat}(Q) \end{array}$$

In substitution it behaves like $\mathbb{B}$ and $\mathbb{U}$ (no surprise since they are each a version of since). It is stable if its first component is stable. It's a local assertion, applying in a single thread, so it can be hatted.

Definition 15: Substitution and hatting of since

$$\begin{array}{rcl} (P\,\mathsf{since}\,Q)[x\backslash\overleftarrow{x}] & = & \overleftarrow{(P\,\mathsf{since}\,Q)} \wedge P[x\backslash\overleftarrow{x}] \\ \overleftarrow{(P\,\mathsf{since}\,Q)} \wedge P & \Rightarrow & (P\,\mathsf{since}\,Q) \\ (P\,\mathsf{since}\,Q)[r\backslash\overleftarrow{r}] & = & P[r\backslash\overleftarrow{r}]\,\mathsf{since}\,Q[r\backslash\overleftarrow{r}] \\ \widehat{P\,\mathsf{since}\,Q} & = & \widehat{P\,\mathsf{since}\,Q} \end{array}$$

By the first two lines of this definition, $r1{=}0 \Rightarrow (\mathbb{U}(x{=}1)\,\mathsf{since}\,y{=}0)$ is stable if $\mathbb{U}(x{=}1)$ is stable, which it is.

Formal reasoning allows us to justify thread postconditions which involve $\mathbb{U}$ and since. But we haven't yet described the machinery which allows us to conclude formally that it is impossible for both threads in fig. 38 to read 0: i.e. that it is impossible at termination that $\mathbb{U}(x{=}1)\,\mathsf{since}\,y{=}0$ could hold in thread 0 and $\mathbb{U}(y{=}1)\,\mathsf{since}\,x{=}0$ in thread 1. The formal machinery is discussed in part V and in our description of the SMT embedding of our logic in appendix B. For the moment we rely on the informal argument we gave above, referring to the thread history diagram of fig. 39.

20 The $\mathbb{S}$ofar modality

As well as $\mathbb{U}$, it is convenient to introduce $\mathbb{S}$ofar. $\mathbb{S}\mathsf{ofar}(P)$ means that P has held since execution started. It has all the properties of definition 4, reading $\mathbb{S}$ofar for $\mathbb{B}$, and of definition 12, reading $\mathbb{S}$ofar for $\mathbb{U}$. In addition we can say that if P has held since the start of execution, then it must certainly have held when interference was propagated or elaborated.

Definition 16: Additional properties of the $\mathbb{S}$ofar modality

$$\begin{array}{rcl} \mathbb{S}\mathsf{ofar}(P) & \Rightarrow & \mathbb{U}(P) \\ \mathbb{S}\mathsf{ofar}(\mathbb{U}(P)) & = & \mathbb{S}\mathsf{ofar}(P) \\ \mathbb{U}(\mathbb{S}\mathsf{ofar}(P)) & = & \mathbb{S}\mathsf{ofar}(P) \\ \mathbb{S}\mathsf{ofar}(P) & \Rightarrow & \widehat{P} \\ \mathbb{S}\mathsf{ofar}(P) & \Rightarrow & \widehat{\widehat{P}} \end{array}$$

In substitution $\mathbb{S}$ofar behaves like since, $\mathbb{U}$ and $\mathbb{B}$; it's unaffected by hatting because (unlike $\mathbb{U}$) $\mathbb{S}\mathsf{ofar}(P) \Rightarrow \widehat{P}$.

Definition 17: Substitution and hatting of the $\mathbb{S}$ofar modality

$$\begin{array}{rcl} \mathbb{S}\mathsf{ofar}(P)[x\backslash\overleftarrow{x}] & = & \overleftarrow{\mathbb{S}\mathsf{ofar}(P)} \wedge P[x\backslash\overleftarrow{x}] \\ \overleftarrow{\mathbb{S}\mathsf{ofar}(P)} \wedge P & \Rightarrow & \mathbb{S}\mathsf{ofar}(P) \\ \mathbb{S}\mathsf{ofar}(P)[r\backslash\overleftarrow{r}] & = & \mathbb{S}\mathsf{ofar}(P[r\backslash\overleftarrow{r}]) \\ \widehat{\mathbb{S}\mathsf{ofar}(P)} & = & \mathbb{S}\mathsf{ofar}(P) \end{array}$$

21 Interference without propagation – UEXT instability

The $\mathbb{U}$ and $\mathbb{S}$ofar modalities introduce a new kind of instability. Suppose we claim, in thread A, that $\mathbb{U}(x{=}1)$. For that claim to be stable, for it to hold no matter how long we wait for thread A to make a move, it must be impossible for any other thread B to write $x{=}0$. We don't have to wait for that write to percolate from B to A: if at any instant $x{=}0$ in thread B, then at that same instant $\neg\mathbb{U}(x{=}1)$ in thread A.

Consider the unproof in fig. 40. The sender is unproblematic, and its guarantee is straightforward – but note that because it uses an lo constraint a→b, rather than bo as in MP, the second line of its guarantee has an amodal interference precondition.

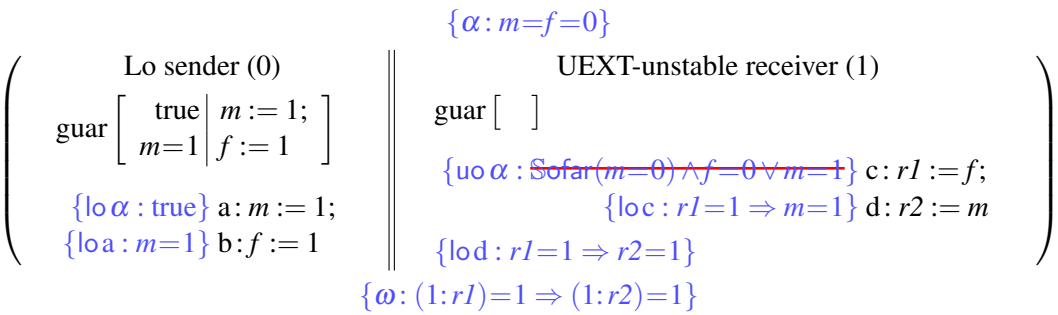
{α : m=f=0}
Lo sender (0)
UEXT-unstable receiver (1)
guar [true | m := 1; ; m=1 | f := 1]
guar [ ]
{lo α : true} a : m := 1;
{lo a : m=1} b : f := 1
{uo α : ~~Sofar(m=0) ∧ f=0 ∨ m=1~~} c : r1 := f;
{loc : r1=1 ⇒ m=1} d : r2 := m
{lo d : r1=1 ⇒ r2=1}
{ω : (1:r1)=1 ⇒ (1:r2)=1}

Figure 40: An unproof of message-passing, displaying UEXT instability

In the receiver the precondition of command c is implied by the initial state: we suppose that our examples are set up so that when each thread starts the initial state has already been propagated to all threads – i.e. $\mathbb{S}\mathsf{ofar}(m{=}f{=}0)$ holds. Then because of the modality properties in definition 4, reading $\mathbb{S}\mathsf{ofar}$ for $\mathbb{B}$, $\mathbb{S}\mathsf{ofar}(m{=}f{=}0)$ implies $\mathbb{S}\mathsf{ofar}(m{=}0) \wedge f{=}0$, which we can weaken to $(\mathbb{S}\mathsf{ofar}(m{=}0) \wedge f{=}0) \vee m{=}1$, the precondition of c.

That precondition is straightforwardly EXT stable against the first line of the sender's guarantee:

$$\begin{array}{ll} & \mathsf{sp}(((\mathbb{S}\mathsf{ofar}(m{=}0) \wedge f{=}0) \vee m{=}1) \wedge \widehat{\mathsf{true}},\ m := 1) \\ = & ((\overleftarrow{\mathbb{S}\mathsf{ofar}(m{=}0)} \wedge \overleftarrow{\widehat{m}}{=}0 \wedge f{=}0) \vee \overleftarrow{\widehat{m}}{=}1) \wedge \mathsf{true} \wedge m{=}1 \\ \Rightarrow & m{=}1 \\ \Rightarrow & (\mathbb{S}\mathsf{ofar}(m{=}0) \wedge f{=}0) \vee m{=}1 \end{array} \tag{35}$$

It's more subtly stable against the second line:

$$\begin{array}{ll} & \mathsf{sp}(((\mathbb{S}\mathsf{ofar}(m{=}0) \wedge f{=}0) \vee m{=}1) \wedge \widehat{m{=}1},\ f := 1) \\ = & ((\mathbb{S}\mathsf{ofar}(m{=}0) \wedge \overleftarrow{f}{=}0) \vee m{=}1) \wedge \widehat{m}{=}1 \wedge f{=}1 \\ \Rightarrow & ((\mathbb{S}\mathsf{ofar}(m{=}0) \wedge \widehat{m}{=}0 \wedge \overleftarrow{f}{=}0) \vee m{=}1) \wedge \widehat{m}{=}1 \wedge f{=}1 \\ \Rightarrow & m{=}1 \\ \Rightarrow & (\mathbb{S}\mathsf{ofar}(m{=}0) \wedge f{=}0) \vee m{=}1 \end{array} \tag{36}$$

– observe that the first implication step uses one of the properties of $\mathbb{S}\mathsf{ofar}$. If $m{=}0$ has held since the beginning of execution then it certainly held when the interference was elaborated – i.e. $\mathbb{S}\mathsf{ofar}(m{=}0) \Rightarrow \widehat{m}{=}0$.

The precondition of d is implied by the postcondition of c:

$$\begin{array}{ll} & \mathsf{sp}((\mathbb{S}\mathsf{ofar}(m{=}0) \wedge f{=}0) \vee m{=}1,\ r1 := f) \\ = & ((\mathbb{S}\mathsf{ofar}(m{=}0) \wedge f{=}0) \vee m{=}1) \wedge r1{=}f \\ \Rightarrow & f{\neq}0 \Rightarrow m{=}1) \wedge r1{=}f \\ \Rightarrow & r1{=}1 \Rightarrow m{=}1 \end{array} \tag{37}$$

and is stable because $m{=}1$ is stable. The rest of the proof is straightforward.

Fig. 40 has to be an unproof: it claims that we can achieve message-passing with an lo-constrained sender. Once again we've overlooked a cause of instability. We can imagine that $f{=}0$ and $m{=}0$ in the receiver, yet $m{=}1$ in the sender, after the first of its interferences has been elaborated but not propagated. At that instant $\mathbb{S}\mathsf{ofar}(m{=}0)$ is already untrue in the receiver since it claims that $m{=}0$ holds in every thread. But m is not yet 1 in the receiver, and so at that instant the precondition of c is false. The EXT stability check supposes that the receiver's precondition holds until the interference is received; our operational intuition says that $\mathbb{U}$-modal and $\mathbb{S}\mathsf{ofar}$-modal assertions may not hold for that long.

This is then an example in which UEXT stability – stability against interference without propagation – really matters. Even though the semantics of $\mathbb{U}$ and $\mathbb{S}\mathsf{ofar}$ have to be global, we can still check UEXT stability locally.

Rule 9: UEXT stability

Constraint embroidery P is UEXT stable against $Q \mid x := E$ from another thread if

$$\mathsf{sp}(\widetilde{P} \wedge Q,\ x := E) \Rightarrow \widetilde{P}$$

where $\widetilde{P}$ replaces both amodal and $\mathbb{B}$-modal occurrences of v with $\widetilde{v}$, leaving only $\mathbb{U}$-modal and $\mathbb{S}\mathsf{ofar}$-modal

occurrences untouched.

Observe that the 'twiddling' in UEXT stability rule 9 is on the opposite side of the conjunction to the 'hatting' in EXT stability rule 7: there we hatted Q because we were checking in the interfered-with thread; here we twiddle P because we are checking in the interfering thread.

The only difference between hatting and twiddling is the treatment of $\mathbb{B}$:

Definition 18: $\mathbb{B}$ and $(\widetilde{\ })$

$$\widetilde{\mathbb{B}(P)} \quad = \quad \mathbb{B}(\widetilde{P})$$

– $\mathbb{B}(P)$ is local, $\mathbb{U}(P)$ and $\mathbb{S}$ofar(P) are not.

The precondition of c in fig. 40 is UEXT unstable against the first line of the sender's guarantee:

$$\begin{array}{rl} & \mathrm{sp}((\mathbb{S}\mathsf{ofar}(m{=}0) \wedge f{=}0) \vee m{=}1)^{\sim} \wedge \mathrm{true},\ m := 1) \\ = & \mathrm{sp}((\mathbb{S}\mathsf{ofar}(m{=}0) \wedge \widetilde{f}{=}0) \vee \widetilde{m}{=}1,\ m := 1) \\ = & ((\overleftarrow{\mathbb{S}\mathsf{ofar}(m{=}0)} \wedge \overleftarrow{\widetilde{m}}{=}0 \wedge \widetilde{f}{=}0) \vee \widetilde{m}{=}1) \wedge m{=}1 \\ \not\Rightarrow & (\mathbb{S}\mathsf{ofar}(m{=}0) \wedge \widetilde{f}{=}0) \vee \widetilde{m}{=}1 \\ = & (\mathbb{S}\mathsf{ofar}(m{=}0) \wedge f{=}0) \vee m{=}1)^{\sim} \end{array} \tag{38}$$

So when we take UEXT stability into account, the precondition of c in the receiver of fig. 40 is unstable and the proof fails, as experiment tells us it should.

22 In-flight UO instability

Uo parallelism is a second kind of in-flight parallelism.

Definition 19: Uo parallelism

If variable assignments are so- but not uo-ordered then the so-later is *uo parallel* with the so-earlier and will interfere with the so-earlier's interference precondition. A variable assignment is also uo parallel with itself.

Note the similarity to bo parallelism (definition 6). One difference is that an assignment is uo parallel with itself: we explore this in the example below. Another is that uo parallelism applies between assignments to the same variable.

Rule 10: UO stability

Interference $P \mid x := E$ is UO stable against uo-parallel $Q \mid y := F$ if

$$\mathrm{sp}(\widetilde{P} \wedge \widetilde{\widetilde{Q}},\ y := F) \Rightarrow \widetilde{P}$$

$\widetilde{\widetilde{Q}}$ is like $\widetilde{Q}$ except that v is replaced by $\widetilde{\widetilde{v}}$ (and, as with a single twiddle, $\mathbb{S}\mathsf{ofar}(P) \Rightarrow \widetilde{\widetilde{P}}$). Note the similarity to BO stability rule 8.

Consider the sender of fig. 41: it's exactly the sender of MP except that we have given command a a $\mathbb{U}$-modal precondition, which is validly inherited from the initial state and is stable because the receiver makes no interference. The sender's guarantee is constructed as usual from the elaboration preconditions. The program works because it has more than the constraints of MP – uo rather than bo, and bo includes uo – but the proof is misleading.

The problem is with the first line of the sender's guarantee. $\mathbb{U}(m{=}0) \mid m := 1$ is a nonsense, because as soon

$$\{\alpha : m{=}f{=}0\}$$

$$\left(\begin{array}{c|c} \text{Uo-unstable Sender (0)} & \text{Receiver (1)} \\ \text{guar}\left[\begin{array}{c|l} \cancel{\mathbb{U}(m{=}0)} & m := 1; \\ \mathbb{B}(m{=}1) & f := 1 \end{array}\right] & \text{guar}\left[\ \right] \\ \{\text{uo}\,\alpha : \mathbb{U}(m{=}0)\}\ \text{a}: m := 1; & \{\text{lo}\,\alpha : f{=}1 \Rightarrow m{=}1\}\ \text{c}: r1 := f; \\ \{\text{boa} : \mathbb{B}(m{=}1)\}\ \text{b}: f := 1 & \{\text{loc} : r1{=}1 \Rightarrow m{=}1\}\ \text{d}: r2 := m \\ & \{\text{lod} : r1{=}1 \Rightarrow r2{=}1\} \end{array}\right)$$

$$\{\omega : (1{:}r1){=}1 \Rightarrow (1{:}r2){=}1\}$$

Figure 41: Unproof which neglects uo self-parallelism of interference

$$\{\alpha : x{=}y{=}0\}$$

$$\left(\begin{array}{c|c} \text{Thread 0} & \text{Thread 1} \\ \{\text{lo}\,\alpha\}\ \text{a}: r1 := y; & \{\text{lo}\,\alpha\}\ \text{c}: r1 := x; \\ \{\text{loa}\}\ \text{b}: x := 1 & \{\text{loc}\}\ \text{d}: y := 1 \end{array}\right)$$

$$\{\omega : \neg((0{:}r1){=}1 \wedge (1{:}r1){=}1)\}$$

Figure 42: LB, laced

as $m := 1$ is elaborated, its interference precondition is false, so we can't say that it will hold when $m{=}1$ is propagated to the receiver. The assignment overtakes *its own interference precondition*.

Rule 10 says that an assignment is uo parallel with itself. In the case of command a the interference precondition isn't UO stable:

$$\begin{array}{ll} & \text{sp}(\widetilde{\mathbb{U}(m{=}0)} \wedge \widetilde{\widetilde{\mathbb{U}(m{=}0)}},\ m := 1) \\ = & \text{sp}(\mathbb{U}(m{=}0) \wedge \widetilde{m}{=}0 \wedge \mathbb{U}(m{=}0) \wedge \widetilde{\widetilde{m}}{=}0,\ m := 1) \\ = & \text{sp}(\mathbb{U}(m{=}0) \wedge \widetilde{m}{=}0 \wedge \widetilde{\widetilde{m}}{=}0,\ m := 1) \\ = & \overleftarrow{\mathbb{U}(m{=}0)} \wedge \overleftarrow{m}{=}0 \wedge \widetilde{m}{=}0 \wedge \widetilde{\widetilde{m}}{=}0 \wedge m{=}1 \\ \not\Rightarrow & \mathbb{U}(m{=}0) \wedge \widetilde{m}{=}0 \\ = & \widetilde{\mathbb{U}(m{=}0)} \end{array} \tag{39}$$

The first line of the guarantee is wrong, and its precondition should be weakened.

It's possible to construct a receiver which shows that relying on the sender's interference as written in fig. 41 is unsound, but it would be tedious to show it.

23 A proof of LB using $\mathbb{S}$ofar

The first of our three weak-memory litmus tests was LB (fig. 4). It was the easiest to explain in operational terms, but it's the hardest of the three to prove, needing heavy temporal machinery despite being apparently so simple. The laced version of LB is in fig. 42. It uses no more than lo ordering in each thread.

The proof of thread 0 is in fig. 43, with a rely from thread 1. It uses auxiliary extension of x and y, as in the proof of WRC (fig. 19), to allow us to distinguish write-1-after-read-0 from write-1-after-read-1. It uses

$$
\begin{array}{l}
\{\alpha : x{=}y{=}(0,0)\} \\
\left(\begin{array}{l}
\qquad\qquad \text{Thread 0} \\
\text{guar}\left[\begin{array}{r|l} \mathbb{B}(x{=}(0,0)) \wedge y{=}(1,0) & x := 1,1; \\ \mathbb{B}(x{=}(0,0)) & x := 1,0 \end{array}\right] \\
\qquad \{\mathsf{lo}\,\alpha : \mathbb{S}\mathsf{ofar}(x{=}(0,0) \wedge (y{=}(0,0) \vee y{=}(1,0)))\}\ \mathrm{a}: r1,\_ := y; \\
\left\{\mathsf{lo}\,\mathrm{a} : \begin{pmatrix} \mathbb{S}\mathsf{ofar}(x{=}(0,0) \wedge (y{=}(0,0) \vee y{=}(1,0))) \\ \wedge\ (r1{=}0 \vee r1{=}1 \wedge y{=}(1,0)) \end{pmatrix}\right\} \\
\qquad\qquad [\mathbb{B}(x{=}(0,0)) \wedge (r1{=}0 \vee r1{=}1 \wedge y{=}(1,0))]\ \mathrm{b}: x := 1, r1 \\
\{\mathsf{lo}\,\mathrm{b} : r1{=}1 \Rightarrow y{=}(1,0) \wedge x{=}(1,1)\} \\
\text{rely}\left[\begin{array}{r|l} \mathbb{B}(y{=}(0,0)) \wedge x{=}(1,0) & y := 1,1; \\ \mathbb{B}(y{=}(0,0)) & y := 1,0 \end{array}\right]
\end{array}\right) \\
\{\omega : \neg((0{:}r1){=}1 \wedge (1{:}r1){=}1)\}
\end{array}
$$

Figure 43: A proof of LB thread 0

$\mathbb{S}$ofar, which allows us to reason about the state in which interference is elaborated and propagated. Because $\mathbb{S}$ofar is susceptible to self-UO instability, command b has an explicit interference precondition.

Our proof is essentially circular. Thread 0 can't see $y{=}(1,1)$ unless it first writes $x{=}(1,0)$ or $x{=}(1,1)$, because reading from one of those writes is the only way that thread 1 can make $r1{=}1$. So before thread 0 writes to x we can be sure that $y{\neq}(1,1)$ – i.e. $y{=}(0,0) \vee y{=}(1,0)$. And then only the second of those possibilities will allow it to write $x{=}(1,1)$. That's captured in its guarantee. And then we know by symmetry that thread 1 must have a similar guarantee, which is thread 1's rely. And that gives us our circular proof: they can't both read after the other writes.

Formally, the precondition of a is EXT stable against the first line of the sender's guarantee because $\mathbb{S}\mathsf{ofar}(x{=}(0,0)) \Rightarrow \widehat{x}{=}(0,0)$:

$$
\begin{array}{lll}
 & \mathsf{sp}(\mathbb{S}\mathsf{ofar}(x{=}(0,0) \wedge (y{=}(0,0) \vee y{=}(1,0))) \wedge (\mathbb{B}(y{=}(0,0)) \wedge x{=}(1,0))\widehat{\ }, y := 1,1) & \\
= & \mathsf{sp}(\mathbb{S}\mathsf{ofar}(x{=}(0,0) \wedge (y{=}(0,0) \vee y{=}(1,0))) \wedge \mathbb{B}(y{=}(0,0)) \wedge \widehat{y}{=}(0,0) \wedge \widehat{x}{=}(1,0), y := 1,1) & \\
\Rightarrow & \mathsf{sp}(\mathbb{S}\mathsf{ofar}(x{=}(0,0) \wedge (y{=}(0,0) \vee y{=}(1,0))) \wedge \widehat{x}{=}(0,0) \wedge \mathbb{B}(y{=}(0,0)) \wedge \widehat{y}{=}(0,0) \wedge \widehat{x}{=}(1,0), y := 1,1) & \\
\Rightarrow & \widehat{x}{=}(0,0) \wedge \widehat{x}{=}(1,0) & \\
= & \text{false} & \\
\Rightarrow & \mathbb{S}\mathsf{ofar}(x{=}(0,0) \wedge (y{=}(0,0) \vee y{=}(1,0))) &
\end{array}
\tag{40}
$$

It's UEXT stable against the same interference because $\mathbb{S}\mathsf{ofar}(x{=}(0,0)) \Rightarrow \widetilde{\widehat{x}}{=}(0,0)$. It's EXT and UEXT stable against the second line of the rely because $y{=}(1,0)$ is stable; it's a register assigment so we don't have to consider BO or UO instability, and it makes no entry in the guarantee.

The postcondition of command a is $\mathbb{S}\mathsf{ofar}(x{=}(0,0) \wedge (y{=}(0,0) \vee y{=}(1,0))) \wedge (r1,\_){=}y$ which implies the precondition of b, which is EXT and UEXT stable against the guarantee for the same reasons as before. But as an interference it would be self-UO unstable so we weaken $\mathbb{S}\mathsf{ofar}(x{=}(0,0) \wedge (y{=}(0,0) \vee y{=}(1,0)))$ to $\mathbb{B}(x{=}(0,0))$ and we have a self-UO-stable interference precondition. Considering the two cases $r1{=}1$ and $r1{=}0$ give us the first and second lines of the guarantee.

The thread postcondition says that if $r1{=}1$ we must have $y{=}(1,0)$; the postcondition of thread 1 would insist that if $r1{=}1$ in thread 1 we must have $y{=}(1,1)$; so we can't have $r1{=}1$ in both threads and we have a

proof.

This is the simplest proof we have been able to construct. It requires temporal cross-thread reasoning, and perhaps that is just the way it has to be. The use of $\mathbb{S}\mathsf{ofar}$ might seem to be restrictive, working only with short and simple litmus tests. But in truth any write can be made unique with just enough auxiliary padding, so though awkward it isn't an incompleteness.

24 The $\mathbb{O}$uat modality

$\mathbb{U}(P)$ says that P has held everywhere since a barrier; $\mathbb{S}\mathsf{ofar}(P)$ that P has held everywhere since execution started. $\mathbb{B}(P)$ says that P has held locally since a barrier. There's clearly a missing local modality.

Rather than make a local analogue of $\mathbb{S}\mathsf{ofar}$, we prefer to make a negated version. $\mathbb{O}\mathsf{uat}(P)$ says that P has held locally sometime since execution started – an 'existence' rather than a 'universal' modality. It's extremely useful when reasoning about the witnessed order of writes to a single location.

Its most spectacular property is that as a positive assertion it's immune to variable interference – if P happened, it happened, and variable interference can't change that. A negative assertion is less stable. Being thread local, it can be hatted and twiddled.

Definition 20: Substitution and hatting of $\mathbb{O}$uat

$$\begin{array}{rcl}
\mathbb{O}\mathsf{uat}(P)[x\backslash\overleftarrow{x}] & = & \overleftarrow{\mathbb{O}\mathsf{uat}(P)} \\
\overleftarrow{\mathbb{O}\mathsf{uat}(P)} \vee P & \Leftrightarrow & \mathbb{O}\mathsf{uat}(P) \\
\neg\overleftarrow{\mathbb{O}\mathsf{uat}(P)} \wedge \neg P & \Leftrightarrow & \neg\mathbb{O}\mathsf{uat}(P) \\
\mathbb{O}\mathsf{uat}(P)[r\backslash\overleftarrow{r}] & = & \mathbb{O}\mathsf{uat}(P[r\backslash\overleftarrow{r}]) \\
\widehat{\mathbb{O}\mathsf{uat}(P)} & = & \widehat{\mathbb{O}\mathsf{uat}(P)}
\end{array}$$

Because it is an existence variety of since rather than a universal, $\mathbb{O}\mathsf{uat}$'s properties are de Morgan opposites of the properties of the other modalities.

Definition 21: Properties of $\mathbb{O}$uat

$$\begin{array}{rcll}
\mathbb{O}\mathsf{uat}(P) & \Leftarrow & P & \\
\mathbb{O}\mathsf{uat}(P \vee Q) & = & \mathbb{O}\mathsf{uat}(P) \vee \mathbb{O}\mathsf{uat}(Q) & \\
\mathbb{O}\mathsf{uat}(\mathbb{O}\mathsf{uat}(P)) & = & \mathbb{O}\mathsf{uat}(P) & \\
\mathbb{O}\mathsf{uat}(P \wedge Q) & \Rightarrow & \mathbb{O}\mathsf{uat}(P) \wedge \mathbb{O}\mathsf{uat}(Q) & \\
\mathbb{O}\mathsf{uat}(P) & \Rightarrow & P & \text{(if } P \text{ doesn't mention variables)} \\
\mathbb{O}\mathsf{uat}(P \wedge Q) & \Leftarrow & \mathbb{O}\mathsf{uat}(P) \wedge \mathbb{O}\mathsf{uat}(Q) & \text{(if } P \text{ or } Q \text{ don't mention variables)} \\
(P \Rightarrow Q) & \Rightarrow & \mathbb{O}\mathsf{uat}(P) \Rightarrow \mathbb{O}\mathsf{uat}(Q) & \text{(if } P \Rightarrow Q \text{ is a tautology)} \\
\mathbb{O}\mathsf{uat}(P) & \Rightarrow & \neg\mathbb{S}\mathsf{ofar}(\neg P) &
\end{array}$$

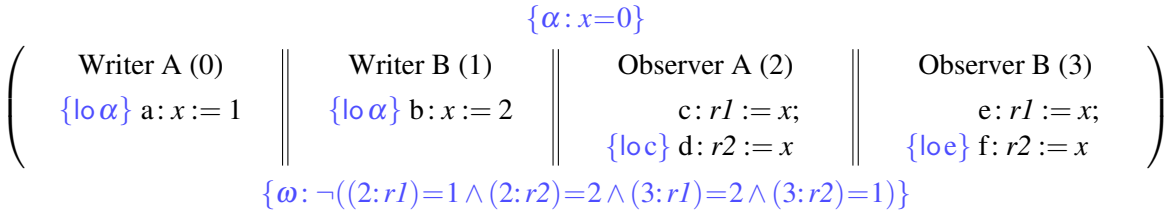

Figure 44: CoRR2 laced

25 Reasoning about coherence order

Several litmus tests and, we believe, many real-world algorithms[13] exploit the 'global coherence' or 'SC per location' constraint SCloc. In those tests and algorithms certain outcomes are impossible because they would require two threads to witness writes to memory in different orders.

Coherence is so important that we have decided to address it explicitly in our logic. Since the constraint applies to writes rather than to values written, we have to ensure that, for variables about which we assert an ordering of writes, each write has a unique value. That involves temporal reasoning and, in many cases, auxiliary information.

25.1 A first coherence example

Consider fig. 44. Writer threads 0 and 1 write different values to x; observer threads 2 and 3 read x twice, each potentially reading from two different writes. Under SCloc each observer thread sees a fraction of the partial order of writes to x, and it is impossible that they could see the same writes in different orders.[14]

The two writers clearly write different values, each different from the initial-state value, and there are no other assignments to x in the program, so we can see that this is an example in which we can identify writes according to the values that they contain.

The proof in fig. 45 contains, in the postcondition of each observer, a coherence assertion $x_c(r1,r2)$, which claims that there are values $W1$ and $W2$ such that $r1{=}W1 \wedge r2{=}W2$, and the write $x{=}W1$ is globally ordered before $x{=}W2$.

For proofchecking purposes we require that every assignment $x := E$ to a variable for which there is a coherence assertion anywhere in the program must have a precondition which implies $\mathbb{S}\mathsf{ofar}(x{\neq}E)$, which guarantees that writes to that variable are unique. This necessarily causes self-UO instability, and to avoid that we shall always need an explicit interference precondition.

Aside 2

There is an argument that $\neg\mathbb{O}\mathsf{uat}(x{=}E)$, which would not provoke self-UO instability, would be sufficient to guarantee uniqueness of writes. Although it's not the case that $\neg\mathbb{O}\mathsf{uat}(x{=}E) \Rightarrow \mathbb{S}\mathsf{ofar}(x{\neq}E)$, we can

[13] One of the origins of this work was Bornat's attempt to verify an idempotent work-stealing algorithm of Michael et al. (2009). That attempt foundered because he couldn't see a way to reason about coherence. Now, four years later, he can't yet reason about the heap. Ho hum, investigations continue.

[14] In the original litmus test this outcome, which involves successive reads from different external writes in each observer thread, would provide implicit dependencies between commands c and d and between e and f. For reasons explained in section 17.1 we must use an explicit lo constraint between the commands in each observer thread, but we believe that nevertheless our treatment captures the essence of the test.

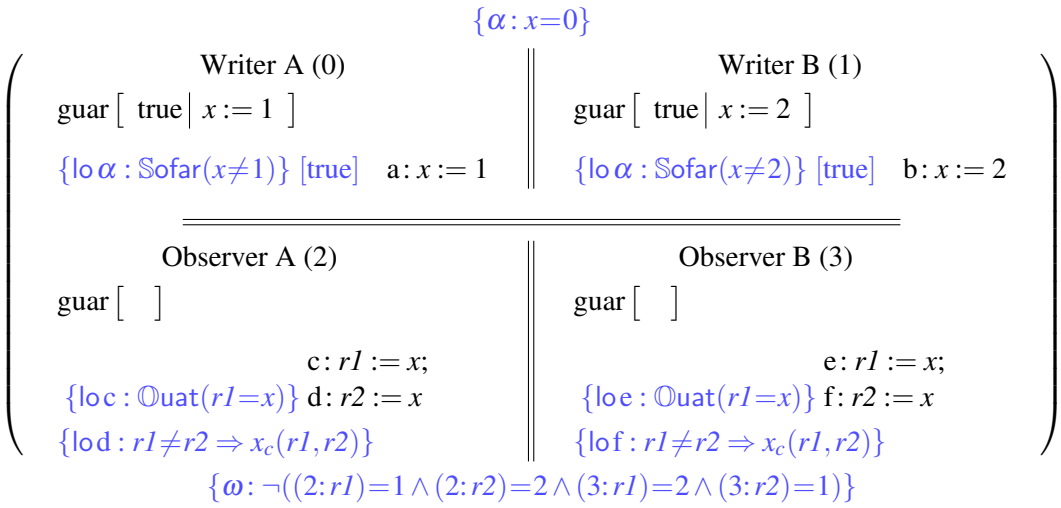
{α : x=0}
Writer A (0)
guar [true | x := 1]
{lo α : 𝕊ofar(x≠1)} [true] a: x := 1
Writer B (1)
guar [true | x := 2]
{lo α : 𝕊ofar(x≠2)} [true] b: x := 2
Observer A (2)
guar [ ]
c: r1 := x;
{loc : ℂuat(r1=x)} d: r2 := x
{lod : r1≠r2 ⇒ xc(r1,r2)}
Observer B (3)
guar [ ]
e: r1 := x;
{loe : ℂuat(r1=x)} f: r2 := x
{lof : r1≠r2 ⇒ xc(r1,r2)}
{ω: ¬((2:r1)=1 ∧ (2:r2)=2 ∧ (3:r1)=2 ∧ (3:r2)=1)}

Figure 45: CoRR2 proof

argue that if there is any assignment $x := E$ that could be executed in any other thread, or earlier in the same thread, then we could never prove $\neg\mathbb{C}\mathsf{uat}(x{=}E)$. In the absence of a soundness proof, this counts as a bonzer wheeze, all of which we've decided to avoid in version 2 of Lace logic. $\mathbb{S}\mathsf{ofar}(x{\neq}E)$ (see especially the embedding description in appendix B) is far more obviously adequate.

Thus the elaboration precondition of command a in the writer asserts $\mathbb{S}\mathsf{ofar}(x{\neq}1)$, the elaboration precondition of b asserts $\mathbb{S}\mathsf{ofar}(x{\neq}2)$, and each command has an interference precondition true. Both elaboration preconditions are implied by the initial state, which establishes $\mathbb{S}\mathsf{ofar}(x{=}0)$. The elaboration precondition of a is stable against the interference of Writer B:

$$\begin{array}{rl} & \mathsf{sp}(\mathbb{S}\mathsf{ofar}(x{\neq}1) \wedge \widehat{\mathsf{true}},\ x := 2) \\ = & \overbrace{\mathbb{S}\mathsf{ofar}(x{\neq}1)} \wedge \overleftarrow{x}{\neq}1 \wedge \mathsf{true} \wedge x{=}2 \\ \Rightarrow & \overbrace{\mathbb{S}\mathsf{ofar}(x{\neq}1)} \wedge x{\neq}1 \\ \Rightarrow & \mathbb{S}\mathsf{ofar}(x{\neq}1) \end{array} \tag{41}$$

and its interference precondition true is trivially implied by $\mathbb{S}\mathsf{ofar}(x{\neq}1)$. The preconditions of b in writer B are valid for similar reasons.

In our treatment coherence is not a modality, it's a relation between values. The meaning of coherence assertions is defined by axioms.

Definition 22: Coherence axioms

membership	$x_c(A,B) \Rightarrow \mathrm{cv}(x)$
irreflexive	$x_c(A,B) \Rightarrow A{\neq}B$
transitive	$x_c(A,B) \wedge x_c(B,C) \Rightarrow x_c(A,C)$
antisymmetric	$x_c(A,B) \Rightarrow \neg x_c(B,A)$
observed	$\mathbb{C}\mathsf{uat}(\mathbb{C}\mathsf{uat}(x{=}A) \wedge x{=}B) \wedge A{\neq}B \wedge \mathrm{cv}(x) \Rightarrow x_c(A,B)$

We require that the unique-write test is applied to variables v for which $v_c(\_,\_)$ appears in any assertion, and implicitly $\mathrm{cv}(v)$ holds only for such variables. The 'membership' axiom then ensures that we can't deduce

$$\{\alpha : m{=}f{=}0\}$$

Thread 0	Thread 1
$\{\mathsf{lo}\,\alpha\}$ a: $m := 1$;	$\{\mathsf{lo}\,\alpha\}$ c: $r1 := f$;
$\{\mathsf{bo}\,\mathrm{a}\}$ b: $f := 1$	$\{\mathsf{lo}\,\mathrm{c};\mathsf{lo}\,\alpha\}$ d: $m := 2$

$$\{\omega : \neg((1{:}r1){=}1 \wedge m{=}1)\}$$

Figure 46: Litmus test S: $\mathbb{B}$ transmits variable history

$x_c(A,B)$ for any other variable.

Coherence is irreflexive, transitive and antisymmetric; $x_c(A,B)$ is observed if in some thread's history there was a state in which $x{=}B$ and previously $x{=}A$; by antisymmetry $x_c(A,B)$ implies that no thread has ever seen $x{=}B$ before $x{=}A$. Because coherence is a global relation, we do not allow deduction of local observations from coherence. In hooking and hatting, only A and B in $x_c(A,B)$ are affected.

Definition 23: Coherence hooking and hatting

$$\begin{array}{rcl}(v_c(A,B))[x\backslash\overleftarrow{x}] & = & v_c(A[x\backslash\overleftarrow{x}],B[x\backslash\overleftarrow{x}])\\ (v_c(A,B))[r\backslash\overleftarrow{r}] & = & v_c(A[r\backslash\overleftarrow{x}],B[x\backslash\overleftarrow{r}])\\ \widehat{v_c(A,B)} & = & v_c(\widehat{A},\widehat{B})\end{array}$$

In observer A of fig. 45 the instantaneous postcondition of c is $r1{=}x$, which is unstable, but which implies the stable $\mathbb{O}\mathsf{uat}(x{=}r1)$. Then the postcondition of d is $\mathbb{O}\mathsf{uat}(x{=}r1) \wedge r2{=}x$ which, provided $r1$ and $r2$ have different values, implies $x_c(r1,r2)$ by the 'observed' axiom of definition 22. Observer B is similar. It is impossible, by the 'antisymmetric' axiom, that we could have $x_c(1,2)$ in one observer and $x_c(2,1)$ in the other, which gives us the program postcondition.

25.2 Coherence and $\mathbb{B}$

Consider fig. 46, a version of the litmus test S.[15] It is almost MP, except that thread 2 doesn't read m, it writes it. But then $m{=}2$ is not stable in thread 0, and we have to reason about coherence (or use auxiliary information) to derive the conclusion.

The proof of S in fig. 47 is straightforward. It contains two coherence assertions for m, so we need to prove $\mathbb{S}\mathsf{ofar}(m{\neq}1)$ in the elaboration precondition of a and $\mathbb{S}\mathsf{ofar}(m{\neq}2)$ in the precondition of d. $\mathbb{S}\mathsf{ofar}(m{\neq}1)$ in thread 0 is stable against true $\mid m := 2$ by an argument very similar to (41), and similarly $\mathbb{S}\mathsf{ofar}(m{\neq}2)$ in thread 1 is stable against true $\mid m := 1$.

The postcondition of a implies $m{=}1$, which implies $\mathbb{O}\mathsf{uat}(m{=}1)$, which is stable. That's decorated with $\mathbb{B}$ to make the embroidery of the a$\rightarrow$b constraint. Each assignment, with its elaboration precondition, is copied to the guarantee.

In thread 1 the preconditions of c and d are stable because $\mathbb{O}\mathsf{uat}(m{=}1)$ is stable and $\mathbb{S}\mathsf{ofar}(m{\neq}2)$ is stable both against true $\mid m := 1$, as noted above, and against $\ldots \mid f := 1$. The postcondition of d implies $(r1{=}1 \Rightarrow \mathbb{O}\mathsf{uat}(m{=}1)) \wedge m{=}2$, which, after some rearrangement and then by the 'observed' axiom, implies the thread postcondition. The program postcondition follows by an argument about eventual consistency which we

[15] The original uses x and y rather than m and f, and swaps 2 and 1. We think our version shows its inheritance from MP more clearly. Note that the initial state needs $m{=}0$ to make sure that writes to m are unique.

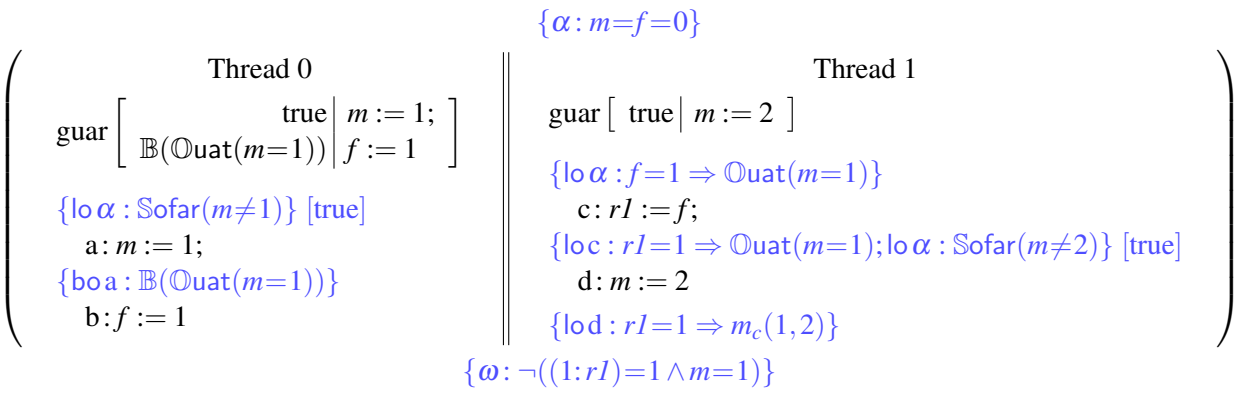
{α : m=f=0}
Thread 0
guar [true | m := 1; 𝔹(𝕆uat(m=1)) | f := 1]
{lo α : 𝕊ofar(m≠1)} [true]
a : m := 1;
{boa : 𝔹(𝕆uat(m=1))}
b : f := 1
Thread 1
guar [true | m := 2]
{lo α : f=1 ⇒ 𝕆uat(m=1)}
c : r1 := f;
{loc : r1=1 ⇒ 𝕆uat(m=1); lo α : 𝕊ofar(m≠2)} [true]
d : m := 2
{lod : r1=1 ⇒ m_c(1,2)}
{ω : ¬((1:r1)=1 ∧ m=1)}

Figure 47: Proof of S

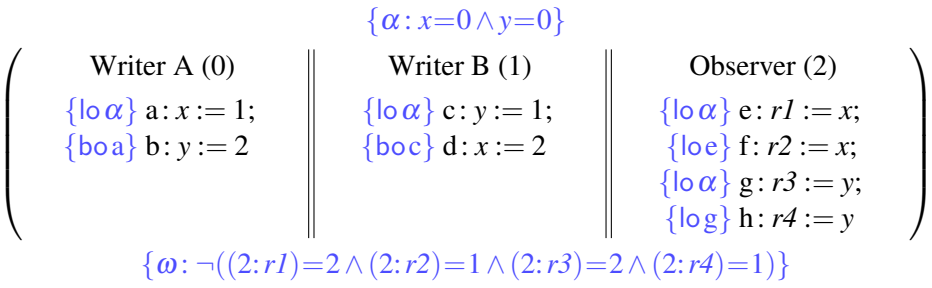

Figure 48: Litmus test 2+2W with observer thread, laced

shall justify in section 29: once the dust has settled, and all the writes have been propagated, $m_c(1,2) \Rightarrow m{\neq}1$.

25.3 2+2W

Fig. 48 is a subtle problem. The online litmus test (Maranget, 2010) has two writer threads with bo ordering in each (but no observer thread) and requires on termination $\neg(x{=}y{=}1)$. The litmus test diagram in (Pow, 2015) requires that with bo ordering we cannot have both $x_c(2,1)$ and $y_c(2,1)$ which, as we shall see in section 29, implies on termination $\neg(x{\neq}2 \wedge y{\neq}2)$. Obviously neither writer can terminate with $x{=}0 \vee y{=}0$. So all of this boils down to $(x{=}1 \wedge y{=}2) \vee (x{=}2 \wedge y{=}1)$. Operational reasoning explains why either requirement must be satisfied: if $x{=}2$ reaches the store before $x{=}1$, for example, then $y{=}1$, which bo-precedes $x{=}2$, must already have arrived and therefore will be before $y{=}2$, which bo-succeeds $x{=}1$ – so on termination $x{=}1$ is the last write to x and $y{=}2$ is the last write to y. And vice-versa if $y{=}1$ is the last write to y, $x{=}2$ will be the last write to x.

The proof in fig. 49 imitates the operational reasoning because we haven't yet found another way. The observer reads x twice in sequence and, internally in parallel, reads y twice. There's a lot of lo-parallel interference but the registers are distinct so there is no lo-parallel instability.

Since the commands which read x and the commands which read y, and their embroidery, are symmetrical, we check the x-reading e and f. The precondition $x{=}2 \Rightarrow \mathbb{O}\mathsf{uat}(y{=}1)$ of e is stable because $\mathbb{O}\mathsf{uat}(y{=}1)$ is stable, and it's non-trivially established by the interference $\mathbb{B}(\mathbb{O}\mathsf{uat}(y{=}1)) \mid x := 2$ of writer B. Then

$$\{\alpha : x{=}0 \wedge y{=}0\}$$

$$\left(\begin{array}{c|c}
\begin{array}{c}\text{Writer A (0)}\\ \text{guar}\left[\begin{array}{l|l}\mathbb{B}(\mathbb{S}\text{ofar}(y{\neq}2)) & x := 1;\\ \mathbb{B}(\mathbb{O}\text{uat}(x{=}1)) & y := 2\end{array}\right]\\ \{\text{lo}\,\alpha : \mathbb{B}(\mathbb{S}\text{ofar}(y{\neq}2))\}\ \text{a} : x := 1;\\ \{\text{bo}\,\text{a} : \mathbb{B}(\mathbb{O}\text{uat}(x{=}1))\}\ \text{b} : y := 2\end{array}
&
\begin{array}{c}\text{Writer B (1)}\\ \text{guar}\left[\begin{array}{l|l}\mathbb{B}(\mathbb{S}\text{ofar}(x{\neq}2)) & y := 1;\\ \mathbb{B}(\mathbb{O}\text{uat}(y{=}1)) & x := 2\end{array}\right]\\ \{\text{lo}\,\alpha : \mathbb{B}(\mathbb{S}\text{ofar}(x{\neq}2))\}\ \text{c} : y := 1;\\ \{\text{boc} : \mathbb{B}(\mathbb{O}\text{uat}(y{=}1))\}\ \text{d} : x := 2\end{array}
\\ \hline
\multicolumn{2}{c}{\begin{array}{c}\text{Observer (2)}\\ \text{guar}\left[\ \right]\\
\{\text{lo}\,\alpha : x{=}2 \Rightarrow \mathbb{O}\text{uat}(y{=}1)\}\ \text{e} : r1 := x;\\
\left\{\text{loe} : r1{=}2 \Rightarrow \left(\begin{array}{l}\mathbb{O}\text{uat}(y{=}1) \wedge \mathbb{O}\text{uat}(x{=}2)\\ \wedge\ (x{=}1 \Rightarrow \neg\mathbb{O}\text{uat}(y{=}1 \wedge \mathbb{O}\text{uat}(y{=}2)))\end{array}\right)\right\}\ \text{f} : r2 := x;\\
\{\text{lo}\,\alpha : y{=}2 \Rightarrow \mathbb{O}\text{uat}(x{=}1)\}\ \text{g} : r3 := y;\\
\left\{\text{log} : r3{=}2 \Rightarrow \left(\begin{array}{l}\mathbb{O}\text{uat}(x{=}1) \wedge \mathbb{O}\text{uat}(y{=}2) \wedge\\ (y{=}1 \Rightarrow \neg\mathbb{O}\text{uat}(x{=}1 \wedge \mathbb{O}\text{uat}(x{=}2)))\end{array}\right)\right\}\ \text{h} : r4 := y\\
\left\{\begin{array}{l}\text{lof} : r1{=}2 \wedge r2{=}1 \Rightarrow \mathbb{O}\text{uat}(x{=}1 \wedge \mathbb{O}\text{uat}(x{=}2)) \wedge \neg\mathbb{O}\text{uat}(y{=}1 \wedge \mathbb{O}\text{uat}(y{=}2));\\ \text{loh} : r3{=}2 \wedge r4{=}1 \Rightarrow \mathbb{O}\text{uat}(y{=}1 \wedge \mathbb{O}\text{uat}(y{=}2)) \wedge \neg\mathbb{O}\text{uat}(x{=}1 \wedge \mathbb{O}\text{uat}(x{=}2))\end{array}\right\}\\
\{\omega : \neg((2{:}r1){=}2 \wedge (2{:}r2){=}1 \wedge (2{:}r3){=}2 \wedge (2{:}r4){=}1)\}\end{array}}
\end{array}\right)$$

Figure 49: 2+2W proof

the postcondition $(x{=}2 \Rightarrow \mathbb{O}\mathsf{uat}(y{=}1)) \wedge r1{=}x$ implies $r1{=}2 \Rightarrow \mathbb{O}\mathsf{uat}(x{=}2)$. Then, we have stability of $x{=}1 \Rightarrow \neg\mathbb{O}\mathsf{uat}(y{=}1 \wedge \mathbb{O}\mathsf{uat}(y{=}2))$ against the interference of writer A:

$$\begin{array}{ll} & \mathrm{sp}((x{=}1 \Rightarrow \neg\mathbb{O}\mathsf{uat}(y{=}1 \wedge \mathbb{O}\mathsf{uat}(y{=}2))) \wedge \mathbb{B}(\widehat{\mathbb{S}\mathsf{ofar}(y{\neq}2)}) \wedge x{=}\widehat{x},\ x := 1) \\ = & (\overleftarrow{x}{=}1 \Rightarrow \neg\mathbb{O}\mathsf{uat}(y{=}1 \wedge \mathbb{O}\mathsf{uat}(y{=}2))) \wedge \mathbb{B}(\mathbb{S}\mathsf{ofar}(y{\neq}2)) \wedge \overleftarrow{x}{=}\widehat{x} \wedge x{=}1 \\ \Rightarrow & \mathbb{S}\mathsf{ofar}(y{\neq}2) \wedge x{=}1 \\ = & \neg\mathbb{O}\mathsf{uat}(y{=}2) \wedge x{=}1 \\ = & \neg\mathbb{O}\mathsf{uat}(\mathbb{O}\mathsf{uat}(y{=}2)) \wedge x{=}1 \\ \Rightarrow & \neg\mathbb{O}\mathsf{uat}(y{=}1 \wedge \mathbb{O}\mathsf{uat}(y{=}2)) \wedge x{=}1 \\ \Rightarrow & x{=}1 \Rightarrow \neg\mathbb{O}\mathsf{uat}(y{=}1 \wedge \mathbb{O}\mathsf{uat}(y{=}2)) \end{array} \tag{42}$$

$$\begin{array}{ll} & \mathrm{sp}((x{=}1 \Rightarrow \neg\mathbb{O}\mathsf{uat}(y{=}1 \wedge \mathbb{O}\mathsf{uat}(y{=}2))) \wedge \mathbb{B}(\widehat{\mathbb{O}\mathsf{uat}(x{=}1)}) \wedge y{=}\widehat{y},\ y := 2) \\ = & \mathrm{sp}((x{=}1 \Rightarrow \overleftarrow{\mathbb{S}\mathsf{ofar}(y{\neq}1 \vee \mathbb{S}\mathsf{ofar}(y{\neq}2))}) \wedge \mathbb{B}(\mathbb{O}\mathsf{uat}(x{=}1)) \wedge y{=}\widehat{y},\ y := 2) \\ \Rightarrow & (x{=}1 \Rightarrow \overleftarrow{\mathbb{S}\mathsf{ofar}(y{\neq}1 \vee \mathbb{S}\mathsf{ofar}(y{\neq}2))}) \wedge (\overleftarrow{y}{\neq}1 \vee \overleftarrow{\mathbb{S}\mathsf{ofar}(y{\neq}2)} \wedge \overleftarrow{y}{\neq}2)) \wedge y{=}2 \\ \Rightarrow & (x{=}1 \Rightarrow \overleftarrow{\mathbb{S}\mathsf{ofar}(y{\neq}1 \vee \mathbb{S}\mathsf{ofar}(y{\neq}2))}) \wedge y{\neq}1 \\ \Rightarrow & (x{=}1 \Rightarrow \mathbb{S}\mathsf{ofar}(y{\neq}1 \vee \mathbb{S}\mathsf{ofar}(y{\neq}2))) \wedge (y{\neq}1 \vee \mathbb{S}\mathsf{ofar}(y{\neq}2)) \\ \Rightarrow & x{=}1 \Rightarrow \mathbb{S}\mathsf{ofar}(y{\neq}1 \vee \mathbb{S}\mathsf{ofar}(y{\neq}2)) \\ = & x{=}1 \Rightarrow \neg\mathbb{O}\mathsf{uat}(y{=}1 \wedge \mathbb{O}\mathsf{uat}(y{=}2)) \end{array} \tag{43}$$

Stability against interference from writer B is trivial, in the first case because we can appeal to $\mathbb{O}\mathsf{uat}(x{=}2)$, and in the other because $x{=}2 \Rightarrow x{\neq}1$

$$\begin{array}{ll} & \mathrm{sp}\left(\begin{pmatrix} \mathbb{O}\mathsf{uat}(x{=}2) \wedge (x{=}1 \Rightarrow \neg\mathbb{O}\mathsf{uat}(y{=}1 \wedge \mathbb{O}\mathsf{uat}(y{=}2))) \\ \wedge\ \mathbb{B}(\widehat{\mathbb{S}\mathsf{ofar}(x{\neq}2)}) \wedge y{=}\widehat{y} \end{pmatrix},\ y := 1\right) \\ \Rightarrow & \mathbb{O}\mathsf{uat}(x{=}2) \wedge \mathbb{B}(\mathbb{S}\mathsf{ofar}(x{\neq}2)) \\ \Rightarrow & \mathbb{O}\mathsf{uat}(x{=}2) \wedge \mathbb{S}\mathsf{ofar}(x{\neq}2) \\ = & \text{false} \\ \Rightarrow & x{=}1 \Rightarrow \neg\mathbb{O}\mathsf{uat}(y{=}1 \wedge \mathbb{O}\mathsf{uat}(y{=}2)) \end{array} \tag{44}$$

$$\begin{array}{ll} & \mathrm{sp}((x{=}1 \Rightarrow \neg\mathbb{O}\mathsf{uat}(y{=}1 \wedge \mathbb{O}\mathsf{uat}(y{=}2))) \wedge \mathbb{B}(\widehat{\mathbb{O}\mathsf{uat}(y{=}1)}) \wedge y{=}\widehat{y},\ x := 2) \\ = & (\overleftarrow{x}{=}1 \Rightarrow \neg\mathbb{O}\mathsf{uat}(y{=}1 \wedge \mathbb{O}\mathsf{uat}(y{=}2))) \wedge \mathbb{B}(\mathbb{O}\mathsf{uat}(y{=}1) \wedge y{=}\widehat{y} \wedge x{=}2 \\ \Rightarrow & x{\neq}1 \\ \Rightarrow & x{=}1 \Rightarrow \neg\mathbb{O}\mathsf{uat}(y{=}1 \wedge \mathbb{O}\mathsf{uat}(y{=}2)) \end{array} \tag{45}$$

Then the postcondition of e can say that if we read $x{=}2$ and then $x{=}1$ we have the makings of $x_c(2,1)$ but we cannot have seen the makings of $y_c(2,1)$ – but note that we can't deduce $\neg y_c(2,1)$ from the coherence axioms of definition 22. The postcondition of h says that if we read $y{=}2$ followed by $y{=}1$ we have the opposite; and we have a contradiction of sorts. It's not the contradiction we wanted, in terms of the values of x and y, but it's sort of convincing: there can't be a thread which reads 2 then 1 from x and also 2 then 1 from y, which is what the litmus test diagram seems to say.

The proof would not have worked with two observer threads. It is entirely possible that one thread can see $x_c(2,1)$ but fail to see $y_c(2,1)$. That doesn't mean that no other thread can see $y_c(2,1)$. Putting the two observer threads into one makes it possible to make a contradiction.

$$\{\alpha : m{=}0 \wedge f{=}0\}$$

$$\left(\begin{array}{c|c|c|c} \text{Sender A (0)} & \text{Sender B (1)} & \text{Receiver A (2)} & \text{Receiver B (3)} \\ \{\mathsf{lo}\,\alpha\}\ \mathrm{a}: m := 1; & \{\mathsf{lo}\,\alpha\}\ \mathrm{a}: m := 2; & \mathrm{e}: r1 := f; & \mathrm{e}: r1 := f; \\ \{\mathsf{boa}\}\ \mathrm{b}: f := 1 & \{\mathsf{boa}\}\ \mathrm{b}: f := 2 & \{\mathsf{loe}\}\ \mathrm{f}: r2 := m & \{\mathsf{loe}\}\ \mathrm{f}: r2 := m \end{array}\right)$$

$$\{\omega : \neg((2{:}r1){=}2 \wedge (2{:}r2){=}1 \wedge (3{:}r1){=}1 \wedge (3{:}r2){=}2)\}$$

Figure 50: A termination puzzle

$$\{\alpha : m{=}0 \wedge f{=}0\}$$

$$\left(\begin{array}{c|c} \text{Sender A (0)} & \text{Sender B (1)} \\ \text{guar}\left[\begin{array}{r|l} \mathbb{B}(\mathbb{S}\mathsf{ofar}(f{\neq}1)) & m := 1; \\ \mathbb{B}(\mathbb{O}\mathsf{uat}(m{=}1)) & f := 1 \end{array}\right] & \text{guar}\left[\begin{array}{r|l} \mathbb{B}(\mathbb{S}\mathsf{ofar}(f{\neq}2)) & m := 2; \\ \mathbb{B}(\mathbb{O}\mathsf{uat}(m{=}2)) & f := 2 \end{array}\right] \\ \{\mathsf{lo}\,\alpha : \mathbb{B}(\mathbb{S}\mathsf{ofar}(f{\neq}1))\}\ \mathrm{a}: m := 1; & \{\mathsf{lo}\,\alpha : \mathbb{B}(\mathbb{S}\mathsf{ofar}(f{\neq}2))\}\ \mathrm{a}: m := 2; \\ \{\mathsf{boa} : \mathbb{B}(\mathbb{O}\mathsf{uat}(m{=}1))\}\ \mathrm{b}: f := 1 & \{\mathsf{boa} : \mathbb{B}(\mathbb{O}\mathsf{uat}(m{=}2))\}\ \mathrm{b}: f := 2 \\ \hline \text{Receiver A (2)} & \text{Receiver B (3)} \\ \text{guar}[\ ] & \text{guar}[\ ] \\ \mathrm{e}: r1 := f; & \mathrm{e}: r1 := f; \\ \{\mathsf{loe} : r1{=}2 \Rightarrow \mathbb{O}\mathsf{uat}(f{=}2)\}\ \mathrm{f}: r2 := m & \{\mathsf{loe} : r1{=}1 \Rightarrow \mathbb{O}\mathsf{uat}(f{=}1)\}\ \mathrm{f}: r2 := m \\ \{\mathsf{lof} : r1{=}2 \wedge r2{=}1 \Rightarrow \mathbb{O}\mathsf{uat}(f{=}2) \wedge m{=}1\} & \{\mathsf{lof} : r1{=}1 \wedge r2{=}2 \Rightarrow \mathbb{O}\mathsf{uat}(f{=}1) \wedge m{=}2\} \end{array}\right)$$

$$\{\omega : \neg((2{:}r1){=}2 \wedge (2{:}r2){=}1 \wedge (3{:}r1){=}1 \wedge (3{:}r2){=}2)\}$$

Figure 51: Apparently-contradictory stable termination states

26 SCloc is not optional

Consider the program in fig. 50: two MPs muddled together. There's a possible execution in which receiver A reads sender B's flag but sender A's message, while receiver B reads A's flag and B's message. The program postcondition says this shouldn't happen. If it did happen we would have, stably, $m{=}1$ in receiver A and $m{=}2$ in receiver B. A proof is shown in fig. 51. The internal deductions are impeccable; this seems to be an impossible outcome.

Under SCloc this outcome really is impossible. When receiver A reads $f{=}2$ we can deduce $\mathbb{O}\mathsf{uat}(m{=}2)$, so when it reads $m{=}1$ we know $m_c(2,1)$. In receiver B, similarly, if we read $f{=}1$ and $m{=}2$ we know $m_c(1,2)$. The global coherence constraint prohibits that outcome, and it would be easy to make a proof to show that it's formally impossible.

But suppose that we don't have SCloc, that there is no global constraint on propagation. Then it would appear that there is an execution of this program which has contradictory stable terminal states of the threads. The informal treatment of termination which we've used up to this point, as well as each of the formal treatments of termination which we present in section 29, would say that the execution is impossible, which would be unsound. But even if we could patch that problem, there's a deeper one.

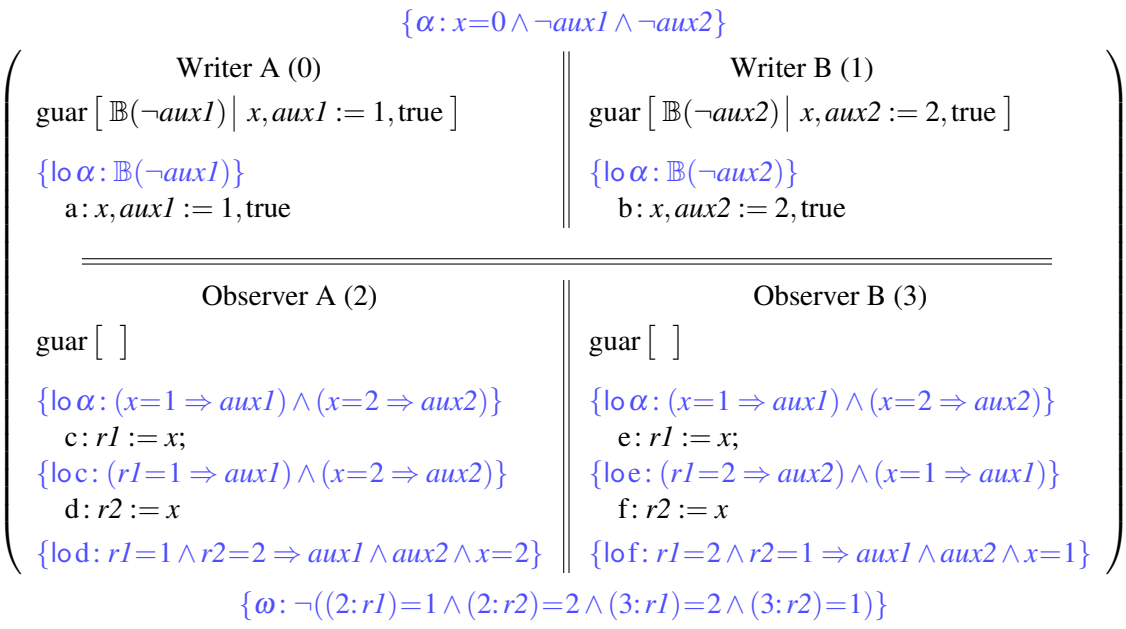
{α : x=0 ∧ ¬aux1 ∧ ¬aux2}
Writer A (0)
guar [𝔹(¬aux1) | x, aux1 := 1, true]
{lo α : 𝔹(¬aux1)}
a : x, aux1 := 1, true
Writer B (1)
guar [𝔹(¬aux2) | x, aux2 := 2, true]
{lo α : 𝔹(¬aux2)}
b : x, aux2 := 2, true
Observer A (2)
guar [ ]
{lo α : (x=1 ⇒ aux1) ∧ (x=2 ⇒ aux2)}
c : r1 := x;
{loc : (r1=1 ⇒ aux1) ∧ (x=2 ⇒ aux2)}
d : r2 := x
{lod : r1=1 ∧ r2=2 ⇒ aux1 ∧ aux2 ∧ x=2}
Observer B (3)
guar [ ]
{lo α : (x=1 ⇒ aux1) ∧ (x=2 ⇒ aux2)}
e : r1 := x;
{loe : (r1=2 ⇒ aux2) ∧ (x=1 ⇒ aux1)}
f : r2 := x
{lof : r1=2 ∧ r2=1 ⇒ aux1 ∧ aux2 ∧ x=1}
{ω : ¬((2:r1)=1 ∧ (2:r2)=2 ∧ (3:r1)=2 ∧ (3:r2)=1)}

Figure 52: CoRR2 with auxiliaries

Our treatment of interference uses assertions to stand for the writes that underpin them. If different threads could receive writes in different orders it is clear that it would not be correct to assume that if P holds stably before a uo barrier then after the barrier P holds in every thread. The writes underpinning P might be propagated everywhere, but some other writes, superseded in P, might be propagated to some of those threads during that process. With SCloc that can't happen; without it we have no valid treatment of $\mathbb{U}$. Similar remarks apply to the use of assertions with $\mathbb{B}$. SCloc is the only mechanism that we are aware of which allows our treatment.

It would seem that we must, reluctantly and in the absence of a better solution, assume SCloc. There is a loss: it means that our logic does not apply to distributed systems, which are very like weak memory but don't have SCloc, or to the non-SCloc GPU architectures which we believe may currently exist.

26.1 Coherence from unrepeatable writes

The proof in fig. 51 relies on unrepeatable writes: once an observer has seen $f=1$ from writer A, for example, we can be sure both that $m=1$ has been propagated to it, and that it won't be propagated again. That's achieved by the the bo ordering in the writer, and enforced by the $\mathbb{B}(\mathbb{S}\mathsf{ofar}(f\neq 1))$ interference precondition in its guarantee.

We can achieve a similar effect in a different example, once again potentially permitting contradictory stable terminal states, using auxiliary information rather than temporal reasoning. We use *composite writes* in which an 'actual write' and an 'auxiliary write' are elaborated simultaneously and propagated simultaneously (see section 31). Fig. 52 is a version of CoRR2; in the final assertion $(2:r1)=1 \wedge (2:r2)=2$ implies that stably $x=2$ in observer A and $(3:r1)=2 \wedge (3:r2)=1$ implies that stably $x=1$ in observer B. Under SCloc, of course, we have $x_c(1,2)$ in observer A and $x_c(2,1)$ in observer B, so it can't happen. Without SCloc it is another example of the possibility of simultaneous contradictory stable states, which we cannot soundly deal with.

Part V

Termination and a loose end

We have based our treatment of interference on the $\mathbb{B}$ modality. A local assertion P is transmuted by a bo-inducing barrier into $\mathbb{B}(P)$, an assertion which guarantees propagation of the writes underlying P before writes generated by later assignments. We were silent about what the 'underlying' or 'underpinning' writes might be for a particular P. The treatment was plausible for atemporal P at least. We used a similar argument when introducing the $\mathbb{U}$ modality: the writes underlying P are transmitted to all threads, and $\mathbb{U}(P)$ claims that P therefore holds in all threads.

Now that we've introduced temporal assertions that argument is no longer sound. Temporal assertions can describe historical coincidences – $\mathbb{O}\mathsf{uat}(x{=}y)$, for example – that can't be said to be described by 'underlying writes', and won't be reliably propagated to other threads. This is a considerable difficulty, which affects our treatment of $\mathbb{B}$, $\mathbb{U}$ and termination.

We resolve the difficulty partly by prohibiting problematic temporal assertions in $\mathbb{B}$ and $\mathbb{U}$, but to treat termination we must grasp the nettle and define what would be propagated by $\mathbb{U}(P)$ if P contains temporal coincidences.

27 $\downarrow P$, the propagatable version of P

Some assertions plainly describe the writes which make up the current state and can obviously be propagated to other threads: $x{=}0$ is an example. Some do not: $\mathbb{O}\mathsf{uat}(x{=}0)$ is an example. But we've used $\mathbb{B}(\mathbb{O}\mathsf{uat}(m{=}1))$ in fig. 47 in the proof of S: were we right to do so? We believe so, because SCloc aka coherence is a feature of our logic. The write(s) that are described by $\mathbb{O}\mathsf{uat}(m{=}1)$ will either include $m{=}1$ or a coherence-successor of it, and that's a useful description. But multivariate historical assertions are a problem: $\mathbb{B}(\mathbb{O}\mathsf{uat}(m{=}1 \wedge f{=}0))$ wouldn't do what it purports to do. By propagating writes from the current state – $m{=}1$ or a coherence-successor, $f{=}0$ or a coherence-successor – it's impossible to guarantee that in the receiving threads' history it was ever the case that $m{=}1 \wedge f{=}0$.

We can, however, define $\downarrow P$, an assertion implied by P that can be sure to be propagated by $\mathbb{B}(P)$ or $\mathbb{U}(P)$. Our definition is incomplete, but that's ok: it just restricts the proofs that we can make. It's complicated by the need to distinguish negative and positive occurrences. In problematic cases we replace the assertion with a fresh boolean variable. $\Downarrow P$ – the negative-position case – is just like $\downarrow P$ except for the modality cases shown.

Definition 24: $\downarrow P$ – positive occurrences – and $\Downarrow P$ – negative occurrences

$$
\begin{array}{rcll}
\downarrow P & = & P & \text{if } P \text{ does not have multiple free variables} \\
\downarrow(\neg P) & = & \neg \Downarrow P & \\
\downarrow(\mathsf{unop}\ P) & = & \mathsf{unop}\ \downarrow P & \\
\downarrow(P\ \mathsf{binop}\ Q) & = & \downarrow P\ \mathsf{binop}\ \downarrow Q & \\
\downarrow(\forall v(P)) & = & \forall v(\downarrow P) & \text{noting that } v \text{ is a bound variable} \\
\downarrow(\exists v(P)) & = & \exists v(\downarrow P) & \text{noting that } v \text{ is a bound variable} \\
\downarrow \mathbb{B}(P) & = & \mathbb{B}(\downarrow P) & \\
\downarrow \mathbb{U}(P) & = & \mathbb{U}(\downarrow P) & \\
\downarrow \mathbb{O}\mathsf{uat}(P \wedge Q) & = & \downarrow \mathbb{O}\mathsf{uat}(P) \wedge \downarrow \mathbb{O}\mathsf{uat}(Q) & \\
\downarrow \mathbb{O}\mathsf{uat}(P \vee Q) & = & \downarrow \mathbb{O}\mathsf{uat}(P) \vee \downarrow \mathbb{O}\mathsf{uat}(Q) & \\
\downarrow \mathbb{O}\mathsf{uat}(P) & = & \mathsf{fresh}\ \mathit{boolv} & \\
\downarrow(P\,\mathsf{since}\,Q) & = & \downarrow P \wedge \downarrow \mathbb{O}\mathsf{uat}(Q) & \\
\downarrow \mathbb{S}\mathsf{ofar}(P \wedge Q) & = & \downarrow \mathbb{S}\mathsf{ofar}(P) \wedge \downarrow \mathbb{S}\mathsf{ofar}(Q) & \\
\downarrow \mathbb{S}\mathsf{ofar}(P) & = & P & \\
\Downarrow \mathbb{B}(P) & = & \mathsf{fresh}\ \mathit{boolv} & \\
\Downarrow \mathbb{U}(P) & = & \mathsf{fresh}\ \mathit{boolv} & \\
\Downarrow \mathbb{O}\mathsf{uat}(P) & = & \mathsf{fresh}\ \mathit{boolv} & \\
\Downarrow(P\,\mathsf{since}\,Q) & = & \mathsf{fresh}\ \mathit{boolv} & \\
\Downarrow \mathbb{S}\mathsf{ofar}(P) & = & \mathsf{fresh}\ \mathit{boolv} &
\end{array}
$$

In computing $\downarrow P$ and $\Downarrow P$ we exploit equivalences, treating $P \Rightarrow Q$ as $\neg P \vee Q$, $P \Leftrightarrow Q$ as $(P \Rightarrow Q) \wedge (Q \Rightarrow P)$, $\neg(P \wedge Q)$ as $\neg P \vee \neg Q$, $\neg(P \vee Q)$ as $\neg P \wedge \neg Q$.

The most important point to note is that in propagation we lose multivariate historical coincidences like $\mathbb{O}\mathsf{uat}(x{=}y)$ or $\mathbb{O}\mathsf{uat}(x{=}1 \wedge \mathbb{O}\mathsf{uat}(y{=}2))$. But single-variable coincidences like $\mathbb{O}\mathsf{uat}(x{=}1 \wedge \mathbb{O}\mathsf{uat}(x{=}2))$ are preserved: they are the province of SCloc.

28 A proper treatment of $\mathbb{B}$ and $\mathbb{U}$

We have a choice in dealing with $\mathbb{B}$- and $\mathbb{U}$-modal temporal assertions. We could say that $\mathbb{B}(P) \Rightarrow \downarrow P$ and $\mathbb{U}(P) \Rightarrow \downarrow P$, and modify our treatment of hooking and hatting to suit, but that would complicate all our definitions. Or we could note that multivariate temporal coincidences in $\mathbb{B}$- or $\mathbb{U}$-modal positions don't mean what they seem to say, and prohibit them. Which is what we have done.

Definition 25: Restricted $\mathbb{B}$, $\mathbb{U}$ and initial assertion

A propagatable assertion P is one for which P is equivalent to $\downarrow P$. An initial-state assertion must be propagatable. In $\mathbb{B}(P)$ and $\mathbb{U}(P)$, P must be propagatable.

Justification for restriction of the initial assertion, and non-restriction of $\mathbb{S}\mathsf{ofar}$, is given in the next section.

$$\{\alpha : x{=}y{=}0\}$$
$$\left(\begin{array}{c|c} \text{Thread 0} & \text{Thread 1} \\ \text{a}: x := 1; & \{\mathsf{lo}\,\alpha\}\ \text{c}: y := 2; \\ \{\mathsf{uo}\,\text{a}; \mathsf{lo}\,\alpha\}\ \text{b}: y := 1 & \{\mathsf{uo}\,\text{c}\}\ \text{d}: r1 := x \end{array}\right)$$
$$\{\omega : \neg((1{:}r1){=}0 \wedge y{=}2)\}$$

Figure 53: Litmus test R: uo ordering essential

29 A proper treatment of termination

The same question of propagation of assertions arises in our treatment of termination. In early examples we pretended that all that was necessary was to take the conjunction of the thread postconditions. In the presence of temporal assertions, that's now obviously inadequate.

We have supposed throughout that our examples are executed within a test harness which ensures that all threads see the initial state before they start, and which somehow brings their final states together when they have all finished.

Initialisation is straightforward: the test harness sets up the variables to satisfy the initial assertion, executes a uo-inducing barrier like MFENCE, sync or dsb, and then sets a flag variable which each thread waits to see. In effect it sets up $\mathbb{S}\mathsf{ofar}(\alpha)$ for the initial assertion α. But since we are relying on it being propagated from the initialising thread to each test-program thread, it must be a propagatable assertion – i.e. no temporal coincidences. But then, the initial state being established, it's possible to note coincidences: so from an initial assertion $x{=}y{=}0$ it's reasonable to deduce $\mathbb{S}\mathsf{ofar}(\mathbb{O}\mathsf{uat}(x{=}y))$ – i.e. $\mathbb{S}\mathsf{ofar}$ assertions don't need to be propagatable.

We can imagine three mechanisms which would deal with termination. They all require that each thread stores the values of its registers into unique shared variables – the $(n{:}reg)$ variables of our final assertions – and signals that it's finished by writing to a thread-specific flag variable. The test harness loops until it sees that all the termination flags have been set and then somehow waits long enough for all the writes of all the threads to have been propagated to it before it inspects the final state. There seem to be three ways to 'wait long enough':

1. The controlling thread might wait a millisecond or so (these machines are very fast);
2. Each test thread might place a bo-inducing barrier before writing its flag;
3. Each test thread might place a uo-inducing barrier before writing its flag.

Because we had no way of deciding which mechanism was best we called the problem the *purple mystic sync* or PMS. Our problem was to devise a logical treatment which fitted any and all of the three mechanisms.

29.1 R: a problematic litmus test

Our difficulties in doing so have to do with temporal assertions. Fig. 53 shows a litmus test which illustrates the challenge. By design on Power and ARM, and experimentally, it is impossible for thread 1 to read 0 from x if globally there is $y_c(1,2)$ coherence and there is a uo-inducing barrier between the commands of each thread. Crucially, the test fails with any weaker constraints – it fails with combinations of a bo-inducing barrier and an lo-inducing address dependency as in fig. 56, for example.

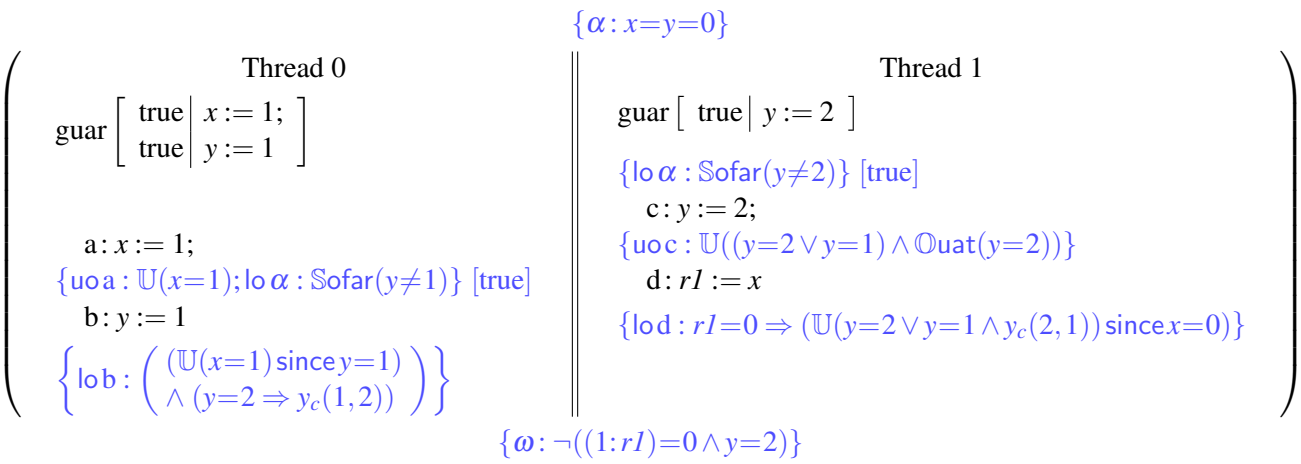
{α : x=y=0}
Thread 0
guar [true | x := 1; true | y := 1]
a : x := 1;
{uo a : 𝕌(x=1); lo α : 𝕊ofar(y≠1)} [true]
b : y := 1
{lo b : ((𝕌(x=1) since y=1) ∧ (y=2 ⇒ y_c(1,2)))}
Thread 1
guar [true | y := 2]
{lo α : 𝕊ofar(y≠2)} [true]
c : y := 2;
{uo c : 𝕌((y=2 ∨ y=1) ∧ 𝕆uat(y=2))}
d : r1 := x
{lo d : r1=0 ⇒ (𝕌(y=2 ∨ y=1 ∧ y_c(2,1)) since x=0)}
{ω : ¬((1:r1)=0 ∧ y=2)}

Figure 54: Proof of R with uo constraints in each thread

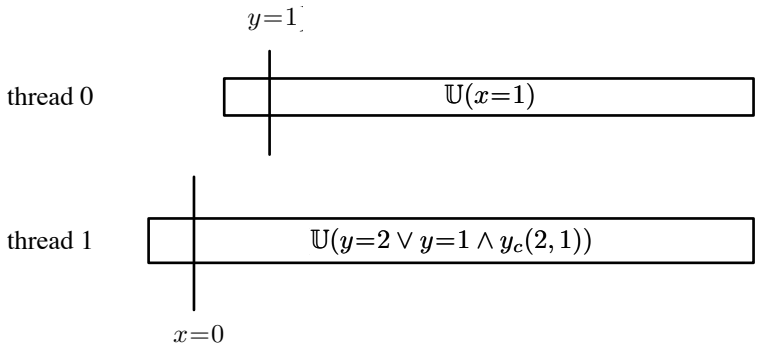

Figure 55: The thread histories of fig. 54

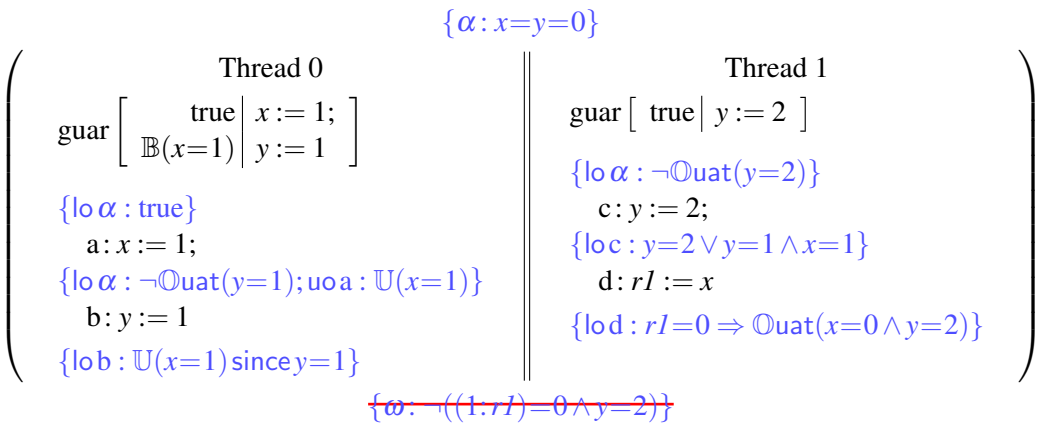
{α : x=y=0}
Thread 0
guar [true | x := 1; ; 𝔹(x=1) | y := 1]
{lo α : true}
a : x := 1;
{lo α : ¬𝕆uat(y=1); uo a : 𝕌(x=1)}
b : y := 1
{lo b : 𝕌(x=1) since y=1}
Thread 1
guar [true | y := 2]
{lo α : ¬𝕆uat(y=2)}
c : y := 2;
{lo c : y=2 ∨ y=1 ∧ x=1}
d : r1 := x
{lo d : r1=0 ⇒ 𝕆uat(x=0 ∧ y=2)}
{ϖ : ¬((1:r1)=0 ∧ y=2)}

Figure 56: An unproof of R with uo in thread 0, lo in thread 1

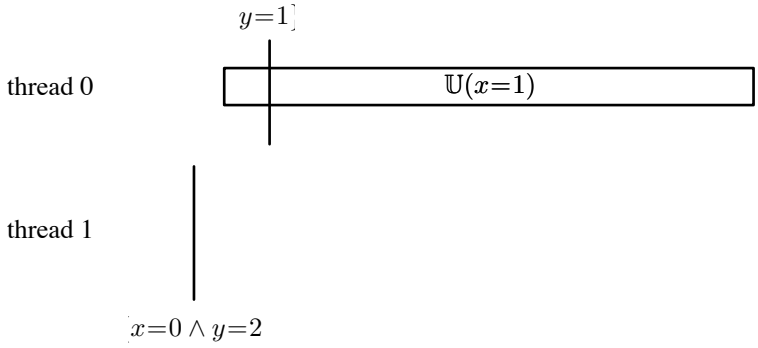

Figure 57: Thread histories of fig. 56

Fig. 54 is a successful proof. The crucial point is that the strongest postcondition of c implies $y{=}2$, which implies $y{=}2 \vee y{=}1 \wedge y_c(2,1)$. That is stable against the interference of thread 0: true $\mid x := 1$ is irrelevant because it doesn't write to y, but $\mathbb{B}(x{=}1) \mid y := 1$ needs checking.

$$
\begin{array}{ll}
 & \mathrm{sp}((y{=}2 \vee y{=}1 \wedge y_c(2,1)) \wedge \widehat{\mathbb{B}(x{=}1)},\ y := 1) \\
= & (\overleftarrow{y}{=}2 \vee \overleftarrow{y}{=}1 \wedge y_c(2,1)) \wedge \mathbb{B}(x{=}1) \wedge \widehat{x}{=}1 \wedge y{=}1 \\
\Rightarrow & (\overleftarrow{y}{=}2 \vee \overleftarrow{y}{=}1 \wedge y_c(2,1)) \wedge y{=}1 \\
= & \overleftarrow{y}{=}2 \wedge y{=}1 \vee \overleftarrow{y}{=}1 \wedge y{=}1 \wedge y_c(2,1) \\
\Rightarrow & \mathbb{O}\mathsf{uat}(y{=}2) \wedge y{=}1 \vee \overleftarrow{y}{=}1 \wedge y{=}1 \wedge y_c(2,1) \\
\Rightarrow & y_c(2,1) \wedge y{=}1 \\
\Rightarrow & y{=}2 \vee y{=}1 \wedge y_c(2,1)
\end{array}
\tag{46}
$$

The universalised version is EXT and UEXT stable, following similar reasoning.

The thread histories are pictured in fig. 55. Plainly $x{=}0$ in thread 1 must precede the $\mathbb{U}(x{=}1)$ period; plainly $\mathbb{U}(y{=}2 \vee y{=}1 \wedge y_c(2,1))$ holds when $y{=}1$ in thread 0, and therefore we have $y_c(2,1)$ from that instant. Plainly? Hardly, and anyway only informally. We need a formal treatment.

To begin to unpick the problem, we present in fig. 56 an unproof of the same test with lo ordering in thread 1. This time the postcondition of c is $y{=}2 \vee y{=}1 \wedge x{=}1$, which is clearly stable. The threads' postconditions describe the thread histories pictured in fig. 57. Plainly $x{=}0 \wedge y{=}2$ in thread 1 occurs globally before the

$\mathbb{U}(x{=}1)$ epoch, which contains $y{=}1$. So $y{=}2$ occurs in thread 1 before $y{=}1$ in thread 0, which seems to suggest that we might have $y_c(2,1)$. Yet experiment shows that we would be wrong to think so, because this program *can* terminate with $(1{:}r1){=}0 \wedge y{=}2$. Plainly we are beyond 'plainly'. Formality is essential.

29.2 The termination problem solved

The conclusion of thread 1 in fig. 56 involves a multivariate temporal coincidence: $\mathbb{O}\mathsf{uat}(x{=}0 \wedge y{=}2)$. None of our PMS mechanisms would propagate this to a controlling thread. Instead it would receive $\downarrow \mathbb{O}\mathsf{uat}(x{=}0 \wedge y{=}2)$, which is $\mathbb{O}\mathsf{uat}(x{=}0) \wedge \mathbb{O}\mathsf{uat}(y{=}2)$, and it couldn't rely on reasoning like fig. 57: the $x{=}0$ event has to come before the $\mathbb{U}(x{=}1)$ epoch, but $y{=}2$ doesn't.

Our treatment of PMS is as follows, where $Q@n$ means 'Q in thread n'.

Definition 26: The PMS assertion

If the final assertions of the threads are $P_0 \ldots P_{n-1}$, in which register occurrences are translated into $(\mathit{threadnum}{:}\mathit{reg})$, and the program-final assertion is ω, then the assertion

$$P_0@0 \wedge P_1@1 \wedge \ldots \wedge P_{n-1}@(n{-}1) \wedge (\downarrow P_0 \wedge \downarrow P_1 \wedge \ldots \wedge \downarrow P_{n-1})@n \Rightarrow \omega@n$$

In checking this assertion we allow an additional *pms coherence* axiom.

Definition 27: PMS coherence axiom

pms coherence $\quad x_c(A,B) \Rightarrow \mathrm{cv}(x) \wedge x{\neq}A$

The PMS assertion allows contradictions between thread-local claims (this is how the termination of SB is handled, for example), but doesn't allow thread-local coincidences to be observed in the controlling thread. The extra coherence axiom allows us to exploit the fact that, when all the program activity is over and writes have been propagated everywhere, all threads see the same last write to each variable.

Thus we check in the proof of fig. 54

$$\left(\begin{array}{l}((\mathbb{U}(x{=}1)\,\mathsf{since}\,y{=}1) \wedge (y{=}2 \Rightarrow y_c(1,2)))@0\\ \wedge\,((1:r1){=}0 \Rightarrow (\mathbb{U}(y{=}2 \vee y{=}1 \wedge y_c(2,1))\,\mathsf{since}\,x{=}0))@1\\ \wedge\,(\mathbb{U}(x{=}1) \wedge \mathbb{O}\mathsf{uat}(y{=}1) \wedge (y{=}2 \Rightarrow y_c(1,2)))@2\\ \wedge\,(((1:r1){=}0 \Rightarrow \mathbb{U}(y{=}2 \vee y{=}1 \wedge y_c(2,1)) \wedge \mathbb{O}\mathsf{uat}(x{=}0)))@2\end{array}\right) \Rightarrow (\neg((1:r1){=}0 \wedge y{=}2))@2 \quad (47)$$

In this example there's a contradiction between the since assertions in threads 1 and 2 in just the case that $(1:r1){=}0 \wedge y{=}2$.

In the case of fig. 56 the PMS assertion is

$$\left(\begin{array}{l}(\mathbb{U}(x{=}1)\,\mathsf{since}\,y{=}1)@0\\ \wedge\,((1{:}r1){=}0 \Rightarrow \mathbb{O}\mathsf{uat}(x{=}0 \wedge y{=}2))@1\\ \wedge\,(\mathbb{U}(x{=}1) \wedge \mathbb{O}\mathsf{uat}(y{=}1))@2\\ \wedge\,(((1{:}r1){=}0 \Rightarrow \mathbb{O}\mathsf{uat}(x{=}0) \wedge \mathbb{O}\mathsf{uat}(y{=}2)))@2\end{array}\right) \Rightarrow (\neg((1{:}r1){=}0 \wedge y{=}2))@2 \quad (48)$$

There's nothing in that assertion which constrains the order of arrival of $y{=}1$, $x{=}0$ and $y{=}2$ in the controlling thread 2. So the assertion doesn't hold, as required.

Part VI
A worked example

```
volatile _Bool latch [WORKERS];
volatile _Bool flag [WORKERS];

void worker(int i) {
  while (!latch[i]);
  for (;;) {
    latch[i] = 0;
    if (flag[i]) {
      flag[i] = 0;
      flag[(i+1)%WORKERS] = 1;
      latch[(i+1)%WORKERS] = 1;
    }
    while (!latch [i]);
  }
}
```

Figure 58: Token ring (from (Alglave et al., 2013))

30 The token ring

The token ring problem was posted on the web as a bug report (wai, 2007), found by Michael Tautschnig, and a model-checked solution was reported in (Alglave et al., 2013). A token (in an element of an array confusingly called 'flag') is passed around a ring of numbered worker threads. Each thread

1. waits until it sees its own element in the array of latches has been set;
2. zeroes that latch element;
3. reads from its own element in the array of tokens;
4. if the token value it sees is non-zero:
   (a) zeroes its own token element;
   (b) and sets the next thread's latch and token.

If a thread sees a zero token then it doesn't pass anything on and circulation stops. A simplification of the original is shown in fig. 58. The bug report observed that the special procedure calls being used to read and set latch and token elements, each of which invoked some barriers, didn't provide all the orderings necessary to avoid a hang-up on a Power machine with eight hardware threads.

In essence each worker is an MP reader, reading latch[i] and flag[i] rather than f and m, and then an MP writer, writing flag[i+1] and latch[i+1] rather than m and f. There's a need to reset latch[i] and flag[i] once they've been read and before they have been written again by another thread, and the time to do that is before writing latch[i+1]. Alglave et al. found bugs by model checking and proposed that those bugs would go away by imposing lo order between the reads of latch[i] and flag[i], via a Power address dependency, and bo between the writes of flag[i+1] and latch[i+1], via a Power lwsync. With those modifications the algorithm model-checked for small sizes of ring, and tested on the hardware that the bug-reporters were using. It seems plausible, given our analysis of MP in the proofs so far, that it should work for a ring of any size running for as long as you like. But weak memory is scary, and it might still bite us. We need a proof that it can't.

Even though it's such a simple program, the token ring is a challenging example. We have solved it in two abstracted guises.

$\{lh{=}tk{=}0\}$
guar $\left[\begin{array}{r|l} [A].A\%W{=}0 \wedge tk{=}lh{=}A & tk := A{+}1; \\ [A].A\%W{=}0 \wedge tk{=}A{+}1 \wedge lh{=}A & lh := A{+}1 \end{array}\right]$
while true do
  $\{lh\%W{=}0 \Rightarrow lh{=}tk\}$
  do $r1 := lh$ until $r1\%W{=}0$;
  $\{r1\%W{=}0 \wedge r1{=}lh{=}tk\}$
  if $r1{=}r2$ then
    $r2 := tk; tk := r2{+}1; lh := r1{+}1$
  fi
od
rely $\left[\begin{array}{r|l} [B].B\%W{\neq}0 \wedge tk{=}lh{=}B & tk := B{+}1; \\ [B].B\%W{\neq}0 \wedge tk{=}B{+}1 \wedge lh{=}B & lh := B{+}1 \end{array}\right]$

Figure 59: Worker 0 of a single-latch, single-token W-ring ($W \geq 2$), SC with assertions

$\{\alpha : lh{=}tk{=}0\}$
Thread 0
x : $rauxN := 0$;
while β : true do
  do
    $\{\mathsf{bo}\,\alpha; \mathsf{lo}\,\mathrm{x}\} \overrightarrow{\vee} (\{\mathsf{lo}\,\mathrm{e}\} \vee \{\mathsf{lo}\,\delta_f\})$ b : $r1 := lh$
  until $\{\mathsf{lo}\,\mathrm{b}\}\gamma$: $r1\%W{=}0$;
  $\{\mathsf{lo}\,\gamma_t\}$ c : $r2 := tk$;
  if $\{\mathsf{lo}\,\mathrm{c}\}\delta$: $r1{=}r2$ then
    $\{\mathsf{lo}\,\delta_t\}$ d : $tk := r2{+}1$;
    $\{\mathsf{lo}\,\mathrm{d}\}$ y : $rauxN := r1{+}W$;
    $\{\mathsf{lo}\,\mathrm{y}; \mathsf{bo}\,\mathrm{d}\}$ e : $lh := r1{+}1$
  fi
od

Figure 60: Worker 0 of a single-latch, single-token W-ring ($W \geq 2$), laced, with auxiliary assignments

$\{\alpha : lh{=}tk{=}0\}$

Thread 0

$$\text{guar}\left[\begin{array}{l|l} [A].A\%W{=}0 \wedge lh{=}A \wedge \mathbb{B}(tk{=}A) & tk := A{+}1; \\ [A].A\%W{=}0 \wedge \mathbb{B}(lh{=}A{=}tk{-}1) & lh := A{+}1 \end{array}\right]$$

$$loopinv \triangleq \left(\begin{array}{l} rauxN{\geq}0 \wedge rauxN\%W{=}0 \\ \wedge\, \exists L(rauxN{-}W{<}L{=}lh{\leq}rauxN \wedge \mathbb{S}\text{ofar}(lh{\leq}rauxN) \wedge \mathbb{B}(L{\leq}tk{\leq}rauxN)) \end{array}\right)$$

$$lhseen \triangleq r1\%W{=}0 \wedge 0{\leq}lh{=}r1{=}rauxN \wedge \mathbb{S}\text{ofar}(lh{\leq}r1) \wedge \mathbb{B}(tk{=}r1)$$

x : $rauxN := 0$;
while β : true do
  do
    $\left(\begin{array}{l} \{\text{bo}\,\alpha : \mathbb{S}\text{ofar}(lh{=}tk{=}0); \text{lo}\,x : rauxN{=}0\} \\ \overrightarrow{\vee}\ (\{\text{lo}\,e : loopinv\} \vee \{\text{lo}\,\delta_f : \text{false}\}) \end{array}\right)$ b : $r1 := lh$
  until $\{\text{lo}\,b : loopinv \wedge (r1\%W{=}0 \Rightarrow r1{=}lh)\}\gamma$: $r1\%W{=}0$;
  $\{\text{lo}\,\gamma_t : lhseen\}$ c : $r2 := tk$;
  if $\{\text{lo}\,c : lhseen \wedge r2{=}tk\}\delta$: $r1{=}r2$ then
    $\{\text{lo}\,\delta_t : lhseen \wedge r1{=}r2{=}tk\}$
    $[r1\%W{=}0 \wedge lh{=}r1{=}r2 \wedge \mathbb{B}(tk{=}r1)]$ d : $tk := r2{+}1$;
    $\{\text{lo}\,d : r1{=}rauxN\}$ y : $rauxN := r1{+}W$;
    $\left\{\begin{array}{l} \text{lo}\,y : rauxN{=}r1{+}W; \\ \text{bo}\,d : \mathbb{B}(r1\%W{=}0 \wedge 0{\leq}lh{=}r1{=}tk{-}1) \wedge \mathbb{S}\text{ofar}(lh{\leq}r1) \end{array}\right\}$
    $[\mathbb{B}(r1\%W{=}0 \wedge 0{\leq}lh{=}r1{=}tk{-}1)]$ e : $lh := r1{+}1$
  fi
od
$$\text{rely}\left[\begin{array}{l|l} [C].C\%W{\neq}0 \wedge lh{=}C \wedge \mathbb{B}(tk{=}C) & tk := C{+}1; \\ [C].C\%W{\neq}0 \wedge \mathbb{B}(lh{=}C{<}tk) & lh := C{+}1 \end{array}\right]$$

Thread 1

$$\text{guar}\left[\begin{array}{l|l} [D].D\%W{=}1 \wedge lh{=}D \wedge \mathbb{B}(tk{=}D) & tk := D{+}1; \\ [D].D\%W{=}1 \wedge \mathbb{B}(lh{=}D{=}tk{-}1) & lh := D{+}1 \end{array}\right]$$

Thread 2

$$\text{guar}\left[\begin{array}{l|l} [E].E\%W{=}2 \wedge lh{=}E \wedge \mathbb{B}(tk{=}E) & tk := E{+}1; \\ [E].E\%W{=}2 \wedge \mathbb{B}(lh{=}E{=}tk{-}1) & lh := E{+}1 \end{array}\right]$$

Figure 61: Proof of worker 0 in a single-latch, single-token W-ring ($W{=}3$)

30.1 A single-latch, single-token ring

One workder from our first abstraction, unlaced as if for SC execution, is in fig. 59; W is the size of the ring. As in C, the '%' operator is mod. In place of arrays of latches and tokens there is a single latch variable *lh* and a single token variable *tk*. Latch and token each carry an ever-increasing integer – in effect, auxiliary information. The assertions are intended to make it clear that this is an MP receiver followed by an MP sender. Since the value of *tk* is never decreased, we don't stop at a zero-value token. Instead our criterion for success in the control expression of the conditional is $r1{=}r2$ – i.e. that *lh* and *tk* are in step. The rely is the union of the guarantees of workers with $B\%W{=}1$, $B\%W{=}2$, …, $B\%W{=}W{-}1$.

The problems of the weak-memory proof are to do with in-flight interference, because the *tk* writes of one worker can overtake the *lh* writes of another, and we have to be careful about bo parallelism between one iteration of the loop and the next.

Fig. 60 shows the single-latch single-token algorithm laced for weak memory. We have added an auxiliary register *rauxN* which records the next *lh*/*tk* value that the worker should deal with. Command b's knot is an iterated constraint, first knot from the initial assertion α and the initialisation of *rauxN*, second knot from the loop-final command e or failure of the control expression δ. Otherwise the lacing is straightforward: lo throughout apart from (as in MP) bo from e to f.

The proof of one worker in a three-worker ring is in fig. 61.[16] We include macros *loopinv* and *lhseen* to shorten the constraint embroidery. Confidence that circulation never stops because latch and token get out of step comes from the fact that δ_f has postcondition false, i.e. the test $r1{=}r2$ always succeeds. Note that, for the first time, this example includes purely auxiliary assignments in x: $rauxN := 0$ and y: $rauxN := r1{+}W$; the constraints x→b and y→e are therefore restricted to not mention regular variables in their embroidery, so that when the auxiliary assignment is deleted, no assertions about regular variable values, which could affect bo or uo propagation, are deleted as well.

The first interesting point is the rely. In SC proofs the rely of one thread is simply the union of the guarantees of the other threads. In weak memory that's not enough, because of bo instability between *tk* and *lh* interference. The *tk* interference of worker 2 bo-interferes with the *lh* interference of worker 1:

$$
\begin{aligned}
&\mathrm{sp}\left(\begin{array}{l}(D\%W{=}1\wedge\mathbb{B}(lh{=}D{=}tk{-}1))\widehat{\ }\wedge(E\%W{=}2\wedge lh{=}E\wedge\mathbb{B}(tk{=}E))\widehat{\widehat{\ }},\\ \quad tk := E{+}1\end{array}\right)\\
&\mathrm{sp}\left(\begin{array}{l}D\%W{=}1\wedge\mathbb{B}(lh{=}D{=}tk{-}1)\wedge\widehat{lh}{=}D{=}\widehat{tk}{-}1\wedge E\%W{=}2\wedge\widehat{\widehat{lh}}{=}E\wedge\mathbb{B}(tk{=}E)\wedge\widehat{\widehat{tk}}{=}E,\\ \quad tk := E{+}1\end{array}\right)\\
=\ &\left(\begin{array}{l}D\%W{=}1\wedge\overleftarrow{\mathbb{B}(lh{=}D{=}tk{-}1)}\wedge lh{=}D{=}\overleftarrow{tk}{-}1\wedge\widehat{lh}{=}D{=}\widehat{tk}{-}1\\ \wedge\ E\%W{=}2\wedge\widehat{\widehat{lh}}{=}E\wedge\overleftarrow{\mathbb{B}(tk{=}E)}\wedge\overleftarrow{tk}{=}E\wedge\widehat{\widehat{tk}}{=}E\\ \wedge\ tk{=}E{+}1\end{array}\right)\\
\Rightarrow\ &D\%W{=}1\wedge E{=}D{+}1\wedge\overleftarrow{\mathbb{B}(lh{=}D{=}tk{-}1)}\wedge lh{=}D\wedge tk{=}E{+}1\\
\not\Rightarrow\ &D\%W{=}1\wedge lh{=}D\wedge\mathbb{B}(lh{=}D{=}tk{-}1)
\end{aligned}
\tag{49}
$$

The same problem would arise in larger rings, and $W{=}3$ is the smallest to show it. We fix the problem by weakening the *lh* interference in the rely to $\mathbb{B}(lh{=}C{<}tk)$.

The initial assertion plus initialisation of *rauxN* establish the loop invariant *loopinv*, which is stable because the rely can only increase *lh* and *tk* until, eventually, $lh\%W{=}0$; the bound on *tk* in *loopinv* means that when

[16] Alone amongst the proofs in this paper, this one has not been completely machine-checked. The SMT solver Z3, which underpins the Arsenic proof checker, can't deal with the modal arithmetic if W is constrained only to be greater than 1, but it is happy if we set W equal to a positive integer of any size we like. That is, we can machine-check the proof for very large values of W, and we are confident anyway that it is a proof, but we don't have the stamp of mechanical approval.

$\{\alpha : lh0{=}1 \wedge lh1{=}lh2{=}lh3{=}0 \wedge tk0{=}1 \wedge tk1{=}tk2{=}tk3{=}0\}$

Thread 0

```
while β : true do
  do
    {lo α} ⃗∨ ({loh} ∨ {lo δ_f ; loc}) b : r1 := lh0
  until {lob} γ : r1 = 1;
   {lo γ_t} c : lh0 := 0;
   {lo γ_t} d : r2 := tk0;
  if {lod} δ : r2 ≠ 0 then
                {lo δ_t} f : tk0 := 0;
                {lo δ_t} g : tk1 := r2 + 1;
     {bof ; bog ; boc} h : lh1 := 1
  fi
od
```

Figure 62: Worker 0 of a multi-latch multi-token four-ring, laced

lh hits that limit, so has tk. The precondition of commands c and d are stable against the rely because $r1\%W{=}0$. Command e is stable against the first line of the rely by the property of $\mathbb{S}$ofar: $\mathbb{S}\mathsf{ofar}(lh{\le}r1) \Rightarrow \widehat{lh}{\le}r1$. The instantaneous postcondition of e implies *loopinv* – trivially, because at that instant $lh\%W{\neq}0 \wedge rauxN\%W{=}0 \wedge rauxN{=}lh{+}W{-}1{=}tk{+}W{-}2$.

There is a little bit of bo parallelism: command e, on one iteration of the loop, is bo-parallel with command d, on the next iteration. But because one demands $\mathbb{B}(lh{=}tk)$ and the other $\mathbb{B}(lh{=}tk{-}1)$, and both require $lh\%W{=}0$, bo stability is straightforward.

Alglave et al. inserted a bo-inducing barrier and an lo-inducing address dependency into the original algorithm. The bo d$\rightarrow$e constraint corresponds to the barrier; the lo b$\rightarrow$c constraint (via γ_t) corresponds to the address dependency. There are a lot of other constraints, but they are all implicitly present in Alglave et al.'s presentation, or can be ignored

- bo $\alpha\rightarrow$b is provided by our supposed test harness;
- lo e$\rightarrow$f is implicit on Power because of SCloc;
- lo $\delta_f\rightarrow$b is embroidered with false;
- lo b$\rightarrow\gamma$, lo c$\rightarrow\delta$ (via γ_t) are implicit on Power because of SCreg;
- constraints to, from and through auxiliary assignments are themselves auxiliary.

So we have added nothing that Alglave et al. did not. It would appear that our constraints are not excessive.

Nevertheless, the single-token single-latch ring is an imperfect abstraction of the original. Although it shows latch and token circulating as they should, it doesn't directly address the zero-token hangup which was the cause of the original bug report, branching on what is in effect auxiliary information in token and latch. And it doesn't show that you can use arrays of latches and tokens to get the right result.

30.2 A multi-latch, multi-token-variable ring

Fig. 62 shows worker 0 from an abstraction with four workers, each with their own latch and token variables: it's a transparent abstraction from fig. 59, if you allow four variables to stand for a four-element array. We

$$notyet(N) \triangleq \mathbb{S}\mathsf{ofar}(xP \leq N \wedge xQ \leq N)$$
$$loopinv \triangleq \begin{pmatrix} rxN\%W = 0 \wedge rxN \geq 0 \wedge rxN - W \leq xQ \leq xP \leq rxN \wedge notyet(rxN) \\ \wedge\ (lh0 = 0 \vee lh0 = 1 \wedge \mathbb{B}(tk0 = rxN + 1 \wedge xP = rxN) \wedge xQ = rxN) \end{pmatrix}$$
$$lhseen \triangleq rxN = xQ \wedge rxN\%W = 0 \wedge rxN \geq 0 \wedge \mathbb{B}(xP = rxN) \wedge notyet(rxN)$$
$$tkseen \triangleq r2 = xQ + 1 > 0 \wedge lhseen$$

Figure 63: token four-ring macros

chose four workers to show that the proof is considerably independent of the size of the ring. The lacing is mostly straightforward, if a little too tight (from h to b) and in one respect a little peculiar (from c, through the 'false' branch of δ to b): we discuss those oddities later.

The proof obligations, on top of those shared with the single-variable abstraction, are to show that it's impossible for spurious repeated interference 'from the past' – from workers which have already contributed to the current state – to reverse the change of *lh0* at command c and *tk0* at command f; and to show that interference 'from the future' can't prematurely undermine the assignment to *tk1* in command g.

Our proof is quite large. We present it in several parts. It depends on auxiliary variables *auxP* and *auxQ* – abbreviated to *xP* and *xQ* to fit the proof on the page – as well as *rauxN* inherited from the single-variable version – abbreviated to *rxN*.

Fig. 64 shows the proof. It's intricate, and there are some new tricks to reveal.

First, the rely is computed from the guarantees of the other three threads, shown in fig. 65. There is some bo instability in this collection: the second line of thread 2's guarantee, which assigns *tk2*:=0, interferes with the precondition of the last line of thread 1's guarantee. But in considering thread 0 we can ignore the effect of assignments to *tk2*, because thread 0 takes no notice of *tk2*. Similarly the second line of thread 3's guarantee interferes with the last line of thread 2's guarantee, and can similarly be ignored.

Second, there is a bo constraint from command z, which is an auxiliary assignment, to command h, which is a regular assignment (albeit in parallel with an auxiliary assignment). Auxiliary-to-regular constraints mustn't add any new orderings to the program: that's ok because everything that z depends on, so does h. The embroidery on an auxiliary-to-regular constraint can't affect the order of propagation: that's ok because the z→h embroidery only mentions auxiliary variables. But note that the barrier between z and h is an auxiliary device just like an auxiliary assignment: in execution terms it isn't really there.

We've introduced a skip command labelled ghbar, which represents the (possibly necessary) barrier between g and h which implements the bo constraint. It's purely an explanatory device.

The proof itself relies mostly on the devices already used in the simpler case of the single-variable token ring. Knowing that *xP* and *xQ* are between $rxN - W$ and *rxN* stops interference from the past; saying that $\mathbb{S}\mathsf{ofar}(xP \leq rxN \wedge xQ \leq rxN)$ stops interference from the future. We also use $\mathbb{B}(xP = rxN)$ in the precondition of the commands up to ghbar: the auxiliary assignment to *xP* has to be bo-ordered with those commands – hence after ghbar – but also bo-ordered before h – hence the auxiliary bo constraint.

We can see that if we ignore assignments to variables which aren't used in the proof of thread 0, all that survives of the interference of thread 2 in the rely of thread 0 are its writes to *xP* and *xQ*, each guarded by a precondition which ensures, by modulo *W* arithmetic, that they can never effectively interfere with anything in thread 0. That makes it clear that we could make a proof for five, six, ... or any number of threads. Not a formal argument, of course, but perhaps good enough.

There's one deficiency in this proof. There's no reason why h should be constrained to elaborate before b

$\{\alpha : lh0{=}1 \wedge lh1{=}lh2{=}lh3{=}0 \wedge tk0{=}1 \wedge tk1{=}tk2{=}tk3{=}0 \wedge xP{=}xQ{=}0\}$

Thread 0

$$\text{guar}\left[\begin{array}{r|l} [A].A\%W{=}0 \wedge A{\geq}0 \wedge \mathbb{B}(xP{=}A) & lh0 := 0; \\ [A].A\%W{=}0 \wedge A{\geq}0 \wedge \mathbb{B}(xP{=}A \wedge tk0{=}A{+}1) \wedge xQ{=}A & tk0 := 0; \\ [A].A\%W{=}0 \wedge A{\geq}0 \wedge \mathbb{B}(xP{=}A) \wedge xQ{=}A & tk1 := A{+}2; \\ [A].A\%W{=}0 \wedge A{\geq}0 \wedge \mathbb{B}(xP{=}A \wedge xQ{=}A) & xP := A{+}1; \\ [A].A\%W{=}0 \wedge A{\geq}0 \wedge \mathbb{B}(xP{=}A{+}1 \wedge xQ{=}A \wedge tk1{=}A{+}2) & lh1, xQ := 1, A{+}1 \end{array}\right]$$

x : $rxN := 0$;
while β : true do
  do
    $\left(\begin{array}{l} \{\text{lo}\,\alpha : \mathbb{S}\text{ofar}(lh0{=}1 \wedge tk0{=}1 \wedge xP{=}xQ{=}0); \text{lo x} : rxN{=}0\} \\ \overrightarrow{\vee}\ (\{\text{lo h} : loopinv\} \vee \{\text{lo}\,\delta_f : \text{false}; \text{lo c} : \text{true}\}) \end{array}\right)$ b : $r1 := lh0$
  until $\{\text{lo b} : loopinv \wedge (r1{=}1 \Rightarrow lhseen \wedge lh0{=}1)\}\gamma$: $r1{=}1$;
  $\{\text{lo}\,\gamma_t : lhseen \wedge lh0{=}1\}\ [rxN\%W{=}0 \wedge rxN{\geq}0 \wedge \mathbb{B}(xP{=}rxN)]$ c : $lh0 := 0$;
  $\{\text{lo}\,\gamma_t : lhseen \wedge \mathbb{B}(tk0{=}xP{+}1)\}$ d : $r2 := tk0$;
  if $\{\text{lo d} : r2{=}rxN{+}1{>}0 \wedge lhseen \wedge \mathbb{B}(tk0{=}r2)\}\delta$: $r2{\neq}0$
  then
    $\{\text{lo}\,\delta_t : r2{=}rxN{+}1{>}0 \wedge lhseen \wedge \mathbb{B}(tk0{=}r2)\}$
    $[rxN\%W{=}0 \wedge rxN{\geq}0 \wedge \mathbb{B}(xP{=}tk0{-}1{=}rxN) \wedge xQ{=}rxN]$ f : $tk0 := 0$;
    $\{\text{lo}\,\delta_t : r2{=}rxN{+}1{>}0 \wedge lhseen\}$
    $[r2{=}rxN{+}1 \wedge rxN\%W{=}0 \wedge rxN{\geq}0 \wedge \mathbb{B}(xP{=}rxN) \wedge xQ{=}rxN]$ g : $tk1 := r2{+}1$;
    $\{\text{lo c} : lh0{=}0; \text{lo f} : r2{=}xP{+}1{>}0 \wedge tk0{=}0 \wedge xP\%W{=}0; \text{lo g} : lhseen\}$ y : $rxN := rxN{+}W$;
    $\left\{\begin{array}{l} \text{lo f} : tk0{=}0 \wedge \mathbb{B}(xP\%W{=}0); \\ \text{lo g} : \exists N(N\%W{=}0 \wedge N{\geq}0 \wedge \mathbb{B}(xP{=}N) \wedge tk1{=}r2{+}1{=}N{+}2 \wedge notyet(N)); \\ \text{lo c} : lh0{=}0 \wedge xQ{=}xP \wedge xP\%W{=}0; \\ \text{lo y} : rxN{=}xP{+}W \wedge xP\%W{=}0 \end{array}\right\}$ ghbar : skip;
    $\left\{\text{bo ghbar} : \left(\begin{array}{l} \exists N(N\%W{=}0 \wedge N{\geq}0 \wedge lh0{=}0 \wedge rxN{=}N{+}W \\ \wedge\ r2{=}N{+}1 \wedge \mathbb{B}(xP{=}N \wedge xQ{=}N) \wedge notyet(N)) \end{array}\right)\right\}$
    $[\exists N(N\%W{=}0 \wedge N{\geq}0 \wedge r2{=}N{+}1 \wedge rxN{=}N{+}W \wedge \mathbb{B}(xP{=}N \wedge xQ{=}N))]$ z : $xP := r2$;
    $\left\{\begin{array}{l} \text{bo z} : \left(\begin{array}{l} rxN\%W{=}0 \wedge rxN{\geq}W \wedge \mathbb{B}(xP{=}rxN{-}W{+}1 \\ \wedge\ xQ{=}rxN{-}W) \wedge \mathbb{S}\text{ofar}(xQ{\leq}rxN{-}W) \end{array}\right); \\ \text{bo ghbar} : \left(\begin{array}{l} lh0{=}0 \wedge \mathbb{B}(tk1{=}r2{+}1 \wedge xQ{=}r2{-}1{=}rxN{-}W) \\ \wedge\ \mathbb{S}\text{ofar}(xP{\leq}rxN{-}W{+}1 \wedge xQ{\leq}rxN{-}W) \wedge rxN\%W{=}0 \end{array}\right) \end{array}\right\}$
    $[r2{=}rxN{-}W{+}1 \wedge rxN\%W{=}0 \wedge rxN{\geq}W \wedge \mathbb{B}(xP{=}xQ{+}1{=}tk1{-}1{=}r2)]$ h : $lh1, xQ := 1, r2$
  fi
od

Figure 64: Proof of worker 0 of a token three-ring (macros in fig. 63, rely in fig. 65)

$$
\left(\begin{array}{c}
\text{Thread 1} \\
\text{guar}\left[\begin{array}{rl}
[B].B\%W=1\wedge B\geq 0\wedge\mathbb{B}(xP=B) & \mid lh1:=0; \\
[B].B\%W=1\wedge B\geq 0\wedge\mathbb{B}(xP=B\wedge tk1=B+1)\wedge xQ=B & \mid tk1:=0; \\
[B].B\%W=1\wedge B\geq 0\wedge\mathbb{B}(xP=B)\wedge xQ=B & \mid tk2:=B+2; \\
[B].B\%W=1\wedge B\geq 0\wedge\mathbb{B}(xP=B\wedge xQ=B) & \mid xP:=B+1; \\
[B].B\%W=1\wedge B\geq 0\wedge\mathbb{B}(xP=B+1\wedge xQ=B\wedge tk2=B+2) & \mid lh2,xQ:=1,B+1
\end{array}\right] \\
\hline\hline
\text{Thread 2} \\
\text{guar}\left[\begin{array}{rl}
[C].C\%W=2\wedge C\geq 0\wedge\mathbb{B}(xP=C) & \mid lh2:=0; \\
[C].C\%W=2\wedge C\geq 0\wedge\mathbb{B}(xP=C\wedge tk2=C+1)\wedge xQ=C & \mid tk2:=0; \\
[C].C\%W=2\wedge C\geq 0\wedge\mathbb{B}(xP=C)\wedge xQ=C & \mid tk3:=C+2; \\
[C].C\%W=2\wedge C\geq 0\wedge\mathbb{B}(xP=C\wedge xQ=C) & \mid xP:=C+1; \\
[C].C\%W=2\wedge C\geq 0\wedge\mathbb{B}(xP=C+1\wedge xQ=C\wedge tk3=C+2) & \mid lh3,xQ:=1,C+1
\end{array}\right] \\
\hline\hline
\text{Thread 3} \\
\text{guar}\left[\begin{array}{rl}
[D].D\%W=3\wedge D\geq 0\wedge\mathbb{B}(xP=D) & \mid lh3:=0; \\
[D].D\%W=3\wedge D\geq 0\wedge\mathbb{B}(xP=D\wedge tk3=D+1)\wedge xQ=D & \mid tk3:=0; \\
[D].D\%W=3\wedge D\geq 0\wedge\mathbb{B}(xP=D)\wedge xQ=D & \mid tk0:=D+2; \\
[D].D\%W=3\wedge D\geq 0\wedge\mathbb{B}(xP=D\wedge xQ=D) & \mid xP:=D+1; \\
[D].D\%W=3\wedge D\geq 0\wedge\mathbb{B}(xP=D+1\wedge xQ=D\wedge tk0=D+2) & \mid lh2,xQ:=1,D+1
\end{array}\right]
\end{array}\right)
$$

Figure 65: Rely of worker 0 of a token four-ring

(on the next iteration). It's clear that b can't read *lh0*=1 until the effect of h's assignment *lh1* := 1 has been propagated to thread 2 and the consequences have filtered right round the ring – so the extra lo ordering used in our proof won't introduce any inefficiency, at least. It's clear also that h has to elaborate before its own next elaboration, so it would be reasonable to lo-constrain it to occur before ghbar rather than b. But to do so would introduce all sorts of lo instability into the proof, requiring still more auxiliary gymnastics. Although we believe such a proof would be possible, we haven't yet attempted to make it.

30.3 Impossible so paths

The single-latch single-token algorithm of fig. 60 and the multi-variable algorithm of fig. 62 both use a conditional δ which is provably certain to deliver true. In the proof of fig. 61 there is a constraint $\delta_f \rightarrow$b which carries embroidery false. There is a path in the so tree which includes δ_f and then b, even though it is provably impossible that no execution could take that path, so the lacing is required.

In the proof of the single-latch single-token ring that spurious constraint has no effect: its embroidery vanishes in the disjunctive precondition of command b. In the multi-latch multi-token proof of fig. 64 things are more complicated, because it is conjoined with an lo constraint c$\rightarrow$b. Without that extra constraint the proof is invalid, because there is an so-tree path c$\rightarrow$d$\rightarrow\delta_f\rightarrow\beta_t\rightarrow$a$\rightarrow$b and, without the lo constraint, the rules of lo parallelism say that c, which zeroes the latch, interferes with and destabilises the precondition of control expression γ on the next iteration of the loop.

We have to include a spurious lo constraint c$\rightarrow$b, to counter lo-parallel instability on an so-tree path that can never be followed. We could argue that no damage has been done to the implementation: SCloc is an assumption of our logic, and so the constraint is perhaps implicit and doesn't require implementation; or we might elide it before we present the program to a compiler, since it is in a knot whose overall assertion is false. But we're not sure we can assume the implicit ordering (ARM is going down some frightening roads), and eliding lacing derived syntactically from a verified proof feels wrong.

We don't believe that we have a solution to the problem of impossible so tree paths and their spurious lacing. At present we acknowledge an incompleteness in our logic in that spurious lacing is sometimes required.

Part VII
And finally

31 Auxiliary variables and registers

Owicki-Gries-style proofs of concurrent algorithms have always needed to use auxiliary state. Our logic is no different. Auxiliary variable names start with *aux*; auxiliary register names start with *raux*. The conditions for sound use of auxiliaries are

1. an expression assigned to an regular register or variable may not mention an auxiliary register;
2. a value from an auxiliary variable may not be assigned to an regular register;
3. a control expression may not mention an auxiliary register;
4. dependencies to or from auxiliary assignments may not induce any additional orderings between regular commands;
5. dependencies to or from auxiliary assignments may not induce any additional propagation orderings of regular variables.

In our notation an expression cannot mention a variable, whether auxiliary or regular, and we cannot assign from a variable to a variable. The first three conditions, then, are versions of the conditions on SC Owicki-Gries proofs, but the fourth and fifth are new for weak memory. Together all five conditions ensure that when auxiliaries are deleted from a program the values assigned to regular registers and variables, the lacing of regular commands and control expressions, and the propagation of regular variable writes will be unchanged.

In our Arsenic proof-checker we check the fourth condition by ensuring that for each auxiliary-to-regular constraint, for each of the immediate regular predecessors of the auxiliary command and for each constrained path between each of those predecessors and the regular target of the constraint, there is a regular path to match each auxiliary path.

We require in our logic that embroidery on all auxiliary-to-regular constraints does not mention regular variables: that ensures the fifth condition and, because of the fourth condition, doesn't introduce any incompleteness.

In weak memory the problem of auxiliaries is propagation. If we assume that assignments to auxiliaries are instantaneously propagated to other threads we can produce unsound proofs. If we put special orderings in our programs to make them propagate before or after other variables, we violate our fourth condition. Our solution is to allow auxiliary writes to 'piggy-back' on regular writes. Auxiliary writes can be made coincidental in elaboration and propagation with a particular regular write and the two can be propagated together as a composite.

We have two logical devices for auxiliary writes. *Auxiliary extension* pretends that the value of an regular variable is a tuple of an regular part and some auxiliary parts where, by convention, the regular part comes first. We allow an extended write and an extended read:

$$x := E, Eaux, Eaux', \ldots \tag{50}$$

$$r, raux, raux', \ldots := x \tag{51}$$

The pretence is that the extended value is propagated as a unit to other threads. This allows simultaneous atomic assignments, but only when involving at most one regular (non-auxiliary) write.

Composite writes allow an regular variable to be assigned at the same time as one or more auxiliaries.

$$x, aux, aux', \cdots := E, Eaux, Eaux', \ldots \tag{52}$$

The writes generated by the assignment are propagated together as a unit. Variables assigned in a composite write may also be assigned in the normal way, and we don't provide a composite read.

32 Inclusion of interference in guarantee (or rely)

Most of the time an assignment corresponds to a single entry in the guarantee, using the assignment precondition as precondition in the entry. But sometimes an assignment corresponds to more than one entry, and sometimes the contribution of several assignments can be combined into a number of entries. We also have to check inclusion of the interferences of other threads when a thread has an explicit rely.

Definition 28: Effect of a single interference

$$\text{intf}(Q \mid x := E) \triangleq Q \wedge x'{=}E \wedge v1{=}v1' \wedge \ldots \wedge vn{=}vn'$$

where $v1,\ldots,vn{=}\overline{vs}\backslash x$

where $\overline{vs}$ are the free variables of the guarantee $\bigcup$ free variables of $Q \bigcup \{x\}$.

Inclusion is tested against the disjunction of the effects of the guarantee, including the base case that nothing changes.

Rule 11: Inclusion of interference

Interference $Q \mid x := E$ is included in the guarantee guar$[g1;\ldots;gn]$ if

$$\text{intf}(Q \mid x := E) \Rightarrow \text{intf}(g1) \vee \ldots \vee \text{intf}(gn) \vee (v1{=}v1' \wedge \ldots \wedge vn{=}vn')$$

where $v1,\ldots,vn{=}\overline{vs}$, the free variables of the guarantee $\bigcup$ free variables of $Q \bigcup \{x\}$.

The same rule applies to an explicit rely and the individual entries in other threads' guarantees: the effect of each entry must be included in the effect of the rely.

33 Library Abstraction and Optimisation

Our logic, if it explains anything, explains weak-memory hardware to programmers. It doesn't solve the problems that face those defining a language such as C11 (C, 2011), and it doesn't solve the problems with the C11 definition uncovered by Batty et al. (2013) because it doesn't deal with library abstraction.

The problem of library abstraction in this context is to define the effects of library code, which operates on shared memory, in such a way that library calls can be embedded in a program that otherwise operates on local memory only, and to be able to reason individually about the effects of such library calls. Our treatment so far doesn't point a way towards library abstraction because the interaction between shared-memory accesses in separate library calls can't be eliminated without heavy uo-inducing barriers. Even though the program outside the library won't cause a problem, the context of other library calls is destabilising, absent those barriers.

Our use of constraints also doesn't sit well with compiler optimisation. We don't know, yet, how optimisations should deal with constraints. The intricacy of our proofs makes it clear that optimisers might have to leave constraints alone unless they have very good meta-level proofs of the properties of their optimisations. That is an argument for library abstraction, of course: if we could isolate library code and constrain optimisations within it, that might be a practicable solution.

The truth is that we have very little to say at this stage about library abstraction and optimisation, and still less to say about how to define programming languages which use one and are subject to the other.

34 Deficiencies

New Lace logic is work in progress. It doesn't yet have a proof of soundness.

Our logic deals only with variables. We hope to move on to heap using separation.

We cannot deal with library abstraction or compiler optimisation.

Our logic can't handle the RDW litmus test.

We have never had more than a sketch of a treatment of sychronising assignment (lwarx / stwcx. on Power, ldrex / strex on ARM, CAS on x86). We haven't yet attempted to handle synchronising assignment in new Lace.

Except for the token ring, we haven't yet verified any non-litmus examples.

Acknowledgements

Thanks to Peter Sewell, Susmit Sarkar, Mike Dodds, Mark Batty, Scott Owens and John Wickerson for extensive discussions when we first envisaged a program logic for Power. Thanks to Derek Williams and Luc Maranget (and most of those above) for putting up with Bornat's ignorant questions about Power.

Thanks to Peter O'Hearn for helping us to begin to make a program logic. Thanks to Noam Riznetsky for telling us about since. Thanks to Nikolaj Bjørner, Josh Berdine and Cristoph Wintersteiger for Z3 help.

Thanks for ever to Crary and Sullivan for pointing the way out of the machine-specificity swamp.

References

Wait latch is vulnerable to weak-memory orderings, August 2007. URL http://www.postgresql.org/message-id/24241.1312739269@sss.pgh.pa.us.

Power ISA Version 2.06. IBM Corp., 2009.

Intel 64 and IA-32 Architectures Software Developer's Manual. Intel Corp., 2009.

ARM Architecture Reference Manual: ARMv7-A and ARMv7-R Edition. ARM Ltd., 2010.

Arsenic proof checker, 2015. URL https://github.com/RBornat/Arsenic. Visited 2015/11/26.

POWER and ARM Litmus Tests, 2015. URL http://www.cl.cam.ac.uk/~pes20/ppc-supplemental/test6.pdf. Visited 2015/08/26.

Parosh Aziz Abdulla, Mohamed Faouzi Atig, Yu-Fang Chen, Carl Leonardsson, and Ahmed Rezine. Counter-Example Guided Fence Insertion under TSO. In *Tools and Algorithms for the Construction and Analysis of Systems*, volume 7214 of *LNCS*, pages 204–219. Springer, 2012.

Parosh Aziz Abdulla, Mohamed Faouzi Atig, Yu-Fang Chen, Carl Leonardsson, and Ahmed Rezine. Memorax, a precise and sound tool for automatic fence insertion under tso. In *TACAS*, pages 530–536. Springer, 2013.

Jade Alglave and Luc Maranget. Stability in Weak Memory Models. In Ganesh Gopalakrishnan and Shaz Qadeer, editors, *Computer Aided Verification*, volume 6806 of *Lecture Notes in Computer Science*, pages 50–66. Springer Berlin / Heidelberg, 2011.

Jade Alglave, Anthony Fox, Samin Ishtiaq, Magnus O. Myreen, Susmit Sarkar, Peter Sewell, and Francesco Zappa Nardelli. The semantics of Power and ARM multiprocessor machine code. In *DAMP 2009: Workshop on Declarative Aspects of Multicore Programming*, January 2009.

Jade Alglave, Luc Maranget, Susmit Sarkar, and Peter Sewell. Litmus: Running tests against hardware. In *Tools and Algorithms for the Construction and Analysis of Systems*, pages 41–44. Springer, 2011.

Jade Alglave, Luc Maranget, Susmit Sarkar, and Peter Sewell. Fences in Weak Memory Models (Extended Version). In *FMSD*, 2012.

Jade Alglave, Daniel Kroening, Vincent Nimal, and Michael Tautschnig. Software verification for weak memory via program transformation. In Matthias Felleisen and Philippa Gardner, editors, *Programming Languages and Systems*, volume 7792 of *Lecture Notes in Computer Science*, pages 512–532. Springer Berlin/Heidelberg, 2013.

Jade Alglave, Luc Maranget, and Michael Tautschnig. Herding cats: Modelling, simulation, testing, and data-mining for weak memory. *TOPLAS*, July 2014. Article No. 7, Issue 2.

Jade Alglave, Mark Batty, Alastair F. Donaldson, Ganesh Gopalakrishnan, Jeroen Ketema, Daniel Poetzl, Tyler Sorensen, and John Wickerson. GPU concurrency: Weak behaviours and programming assumptions. In *ASPLOS*, 2015.

Mohamed Atig, Ahmed Bouajjani, and Gennaro Parlato. Getting Rid of Store-Buffers in TSO Analysis. In Ganesh Gopalakrishnan and Shaz Qadeer, editors, *Computer Aided Verification*, volume 6806 of *Lecture Notes in Computer Science*, pages 99–115. Springer Berlin / Heidelberg, 2011.

Mohamed Faouzi Atig, Ahmed Bouajjani, Sebastian Burckhardt, and Madanlal Musuvathi. On the verification problem for weak memory models. In *Proceedings of the 37th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, POPL '10, pages 7–18, 2010.

Mohamed Faouzi Atig, Ahmed Bouajjani, Sebastian Burckhardt, and Madanlal Musuvathi. What's Decidable about Weak Memory Models? In *Programming Languages and Systems*, volume 7211 of *LNCS*, pages 26–46. Springer, 2012.

Mark Batty, Scott Owens, Susmit Sarkar, Peter Sewell, and Tjark Weber. Mathematizing C++ concurrency. In *POPL*, pages 55–66. ACM, 2011.

Mark Batty, Mike Dodds, and Alexey Gotsman. Library abstraction for C/C++ concurrency. In *POPL*, pages 235–248. ACM, 2013.

Hans-J. Boehm and Sarita V. Adve. You Don't Know Jack About Shared Variables or Memory Models. *Communications of the ACM*, 55(2):48–54, February 2012. doi: 10.1145/2076450.2076465.

Hans-Juergen Boehm and Sarita V. Adve. Foundations of the C++ concurrency memory model. In *PLDI*, pages 68–78. ACM, 2008.

Hans-Juergen Boehm and Brian Demsky. Outlawing ghosts: avoiding out-of-thin-air results. In *MSPC*, 2014.

Sebastian Burckhardt and Madan Musuvathi. Effective Program Verification for Relaxed Memory Models. In *CAV*, 2008.

Sebastian Burckhardt, Rajeev Alur, and Milo M. K. Martin. Checkfence: checking consistency of concurrent data types on relaxed memory models. In *Proceedings of the 2007 ACM SIGPLAN conference on Programming language design and implementation*, PLDI '07, pages 12–21, 2007.

C. ISO SC22/WG14. ISO/IEC 9899: 2011 — Programming language C, 2011.

C++. ISO International Standard ISO/IEC 14882:2014(E) — Programming Language C++, 2014. URL `isocpp.org/std/the-standard`. Visited 2015/10/20.

Cristiano Calcagno, Peter O'Hearn, and Hongseok Yang. Concurrent Abstract Separation Logic. in preparation, 2006.

Karl Crary and Michael J. Sullivan. A calculus for relaxed memory. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '15, pages 623–636, New York, NY, USA, 2015. ACM. ISBN 978-1-4503-3300-9. doi: 10.1145/2676726.2676984. URL `http://doi.acm.org/10.1145/2676726.2676984`.

Leonardo De Moura and Nikolaj Bjørner. Z3: An efficient smt solver. In *Tools and Algorithms for the Construction and Analysis of Systems*, pages 337–340. Springer, 2008.

Edsger W. Dijkstra. *A Discipline of Programming*. Prentice-Hall, 1976.

Edsger W Dijkstra. About the sequentiality of process descriptions, n.d. URL `http://www.cs.utexas.edu/users/EWD/translations/EWD35-English.html`. Visited 2015/08/28.

Thomas Dinsdale-Young, Mike Dodds, Philippa Gardner, Matthew J Parkinson, and Viktor Vafeiadis. Concurrent Abstract Predicates. In *ECOOP 2010–Object-Oriented Programming*, pages 504–528. Springer, 2010.

Robert W. Floyd. Assigning Meanings to Programs. In J. T. Schwartz, editor, *Mathematical Aspects of Computer Science*, volume 19 of *Proceedings of Symposia in Applied Mathematics*, pages 19–32, Providence, Rhode Island, 1967. American Mathematical Society.

James Gosling, Bill Joy, Guy Steele, Gilad Bracha, and Alex Buckley. The Java® Language Specification, Java SE 8 Edition, 2015. URL `https://docs.oracle.com/javase/specs/jls/se8/jls8.pdf`. Visited 2015/10/20.

C. A. R. Hoare. An axiomatic basis for computer programming. *Commun. ACM*, 12(10):576–580, 1969. ISSN 0001-0782. doi: http://doi.acm.org/10.1145/363235.363259.

Thuan Huynh and Abhik Roychoudhury. Memory model sensitive bytecode verification. *Formal Methods in System Design*, 31:281–305, 2007.

Cliff B. Jones. Specification and design of (parallel) programs. In *Proceedings of IFIP '83*, pages 321–332, 1983.

M Kuperstein, Martin T. Vechev, and E Yahav. Automatic inference of memory fences. In *Formal Methods in Computer-Aided Design (FMCAD)*, pages 111–119, 2010.

Ori Lahav and Viktor Vafeiadis. Owicki-Gries Reasoning for Weak Memory Models. In *ICALP*, 2015.

Leslie Lamport. How to make a multiprocessor computer that correctly executes multiprocess programs. *IEEE Transactions on Computers*, 100(9):690–691, 1979.

A. Linden and P.Wolper. A verification-based approach to memory fence insertion in relaxed memory systems. In *SPIN*, 2011.

Sela Mador-Haim, Luc Maranget, Susmit Sarkar, Kayvan Memarian, Jade Alglave, Scott Owens, Rajeev Alur, Milo M. K. Martin, Peter Sewell, and Derek Williams. An axiomatic memory model for POWER multiprocessors. In *CAV*, pages 495–512. Springer, 2012.

Luc Maranget. 2+2W+lwsyncs, 2010. URL `http://moscova.inria.fr/~maranget/AXIOM/2+2W+lwsyncs.html`. visited 2016/09/21.

Luc Maranget, Susmit Sarkar, and Peter Sewell. A Tutorial Introduction to the ARM and POWER Relaxed Memory Models, 2012. URL `http://www.cl.cam.ac.uk/~pes20/ppc-supplemental/test7.pdf`. Visited 2015/09/04.

Paul E. McKenney. *Exploiting Deferred Destruction: An Analysis of Read-Copy-Update Techniques in Operating System Kernels*. PhD thesis, OGI School of Science and Engineering, Oregon Health and Sciences University, 2004. URL `http://www.rdrop.com/users/paulmck/RCU/RCUdissertation.2004.07.14e1.pdf`.

Maged M. Michael, Martin T. Vechev, and Vijay A. Saraswat. Idempotent work stealing. In *Proceedings of the 14th ACM SIGPLAN symposium on Principles and practice of parallel programming*, pages 45–54, 2009.

Brian Norris and Brian Demsky. Cdschecker: checking concurrent data structures written with C/C++ atomics. In *OOPSLA*, 2013.

Scott Owens, Susmit Sarkar, and Peter Sewell. A better x86 memory model: x86-TSO. In *TPHOLs*, pages 391–407. Springer, 2009.

S. Owicki and D. Gries. An axiomatic proof technique for parallel programs. *Acta Informatica*, 19:319–340, 1976a.

S. Owicki and D. Gries. Verifying Properties of Parallel Programs: an axiomatic approach. *Commun. ACM*, 19(5):279–285, 1976b.

Susan Speer Owicki. *Axiomatic proof techniques for parallel programs.* PhD thesis, Cornell, 1975. Technical report TR75-251.

John C. Reynolds. Separation Logic: A Logic for Shared Mutable Data Structures. In *LICS '02: Proceedings of the 17th Annual IEEE Symposium on Logic in Computer Science*, pages 55–74, 2002.

Tom Ridge. A rely-guarantee proof system for x86-TSO. In *Verified Software: Theories, Tools, Experiments*, pages 55–70. Springer, 2010.

Susmit Sarkar, Peter Sewell, Francesco Zappa Nardelli, Scott Owens, Tom Ridge, Thomas Braibant, Magnus O. Myreen, and Jade Alglave. The semantics of x86-CC multiprocessor machine code. In *POPL*, pages 379–391. ACM, 2009.

Susmit Sarkar, Peter Sewell, Jade Alglave, Luc Maranget, and Derek Williams. Understanding POWER multiprocessors. In *Proceedings of the 32nd ACM SIGPLAN conference on Programming language design and implementation*, pages 175–186, 2011.

Susmit Sarkar, Kayvan Memarian, Scott Owens, Mark Batty, Peter Sewell, Luc Maranget, Jade Alglave, and Derek Williams. Synchronising C/C++ and POWER. In *PLDI*, pages 311–322. ACM, 2012.

Jaroslav Sevcík, Viktor Vafeiadis, Francesco Zappa Nardelli, Suresh Jagannathan, and Peter Sewell. Relaxed-memory concurrency and verified compilation. In *POPL*, 2011.

Peter Sewell, Susmit Sarkar, Scott Owens, Francesco Zappa Nardelli, and Magnus O. Myreen. x86-TSO: A rigorous and usable programmer's model for x86 multiprocessors. *Communications of the ACM*, 53(7): 89–97, July 2010. (Research Highlights).

Filip Sieczkowski. *Towards Modular Reasoning for Realistic Programming Languages*. PhD thesis, IT University of Copenhagen, November 2013.

Filip Sieczkowski, Kasper Svendsen, Lars Birkedal, and Jean Pichon-Pharabod. A separation logic for fictional sequential consistency. In *ESOP*, 2015.

Joseph Tassarotti, Derek Dreyer, and Viktor Vafeiadis. Verifying read-copy-update in a logic for weak memory. In *PLDI*, 2015.

Aaron Turon, Viktor Vafeiadis, and Derek Dreyer. GPS: Navigating weak memory with ghosts, protocols, and separation. In *Proceedings of the 2014 ACM International Conference on Object Oriented Programming Systems Languages & Applications*, pages 691–707. ACM, 2014.

Viktor Vafeiadis. *Modular fine-grained concurrency verification*. PhD thesis, University of Cambridge, 2007.

Viktor Vafeiadis. Formal reasoning about the C11 weak memory model. In *Proceedings of the 2015 Workshop on Certified Programs and Proofs*, pages 1–2. ACM, 2015.

Viktor Vafeiadis and Chinmay Narayan. Relaxed separation logic: A program logic for c11 concurrency. In *ACM SIGPLAN Notices*, volume 48, pages 867–884. ACM, 2013.

Viktor Vafeiadis and Matthew J. Parkinson. A Marriage of Rely/Guarantee and Separation Logic. In *CONCUR 2007 – Concurrency Theory*, volume 4037 of *LNCS*, pages 256–271, August 2007.

Viktor Vafeiadis and Francesco Zappa Nardelli. Verifying fence elimination optimisations. In Eran Yahav, editor, *Static Analysis*, volume 6887 of *Lecture Notes in Computer Science*, pages 146–162. Springer Berlin / Heidelberg, 2011.

Viktor Vafeiadis, Thibaut Balabonski, Soham Chakraborty, Robin Morisset, and Francesco Zappa Nardelli. Common Compiler Optimisations are Invalid in the C11 Memory Model and what we can do about it. In *POPL*, 2015.

Ian Wehrman and Josh Berdine. A proposal for weak-memory local reasoning. In *LOLA workshop*, 2011.

Qiwen Xu, Willem-Paul de Roever, and Jifeng He. The rely-guarantee method for verifying shared variable concurrent programs. *Formal Aspects of Computing*, 9(2):149–174, 1997.

A Summary of the logic

The notation is described in table 2 and table 1. In text input for the Arsenic proof checker, *knot* brackets are "{⋆" and "⋆}", and *intfpre* brackets are "[⋆" and "⋆]".

Auxiliary extension of variables, restrictions on auxiliary assignment and auxiliary-to-regular lacing are described in section 31 and are not summarised here.

Constraints run between labelled primitive components – assignments and control expressions. Constraint lacing may only reinforce the dependencies of the so tree. Where a component occurs more than once in the so tree, it may need a disjunctive constraint constructed according to definition 7. Components in loops may require iterated constraints.

If a component is constrained at all, then for all paths in the so tree to an instance of that component, there must be a disjunct in the component's knot which depends on one of the nodes of that path.

A component x in a loop may have an iterated constraint *knot1* $\overrightarrow{\vee}$ *knot2*. Constraints in *knot1* must cover all paths from the initial assertion to x; constraints in *knot2* must cover all paths from x to x.

Interference and elaboration preconditions are derived from knots according to definition 11.

- The overall precondition is the conjunction of a knot's embroidery or, if the knot is a disjunction of knots, the disjunction of their overall preconditions.
- The elaboration precondition is the overall precondition, excluding the embroidery of go constraints, conjoined with the satisfaction (sat) of the overall precondition.
- The interference precondition is either the overall precondition or, if the knot contains a square-bracketed *intfpre* assertion (see table 2), the *intfpre* assertion, which must be implied by the overall precondition.

The embroidery of a constraint is derived from the postcondition of its source component. Component postconditions are derived according to definitions 1 and 8. Strongest postconditions are defined in definition 2.

The postcondition of the initial-state assertion is the assertion itself. The postcondition of skip is its precondition. The postcondition of assert P is the conjunction of its precondition and P. For assignments we use *strongest postconditions* (Dijkstra, 1976).

A control expression γ: E with elaboration precondition P has two postconditions, one for each of its outcomes. The γ_t outcome has $P \wedge E$; the γ_f outcome has $P \wedge \neg E$.

$$\begin{array}{lcl} \mathrm{sp}(P,\ x := E) & \triangleq & P[x\backslash\overleftarrow{x}] \wedge x = E \\ \mathrm{sp}(P,\ r := E) & \triangleq & P[r\backslash\overleftarrow{r}] \wedge r = E[r\backslash\overleftarrow{r}] \\ \mathrm{sp}(P,\ r := x) & \triangleq & P[r\backslash\overleftarrow{r}] \wedge r = x \end{array}$$

Assertions embroidered on stitches must be inherited as described in rule 2.

If a constraint's source postcondition is P and the constraint's embroidery is Q then we require

- on an lo or go constraint, $P \Rightarrow Q$;
- on a bo constraint, for some R, $P \Rightarrow R$ and $\mathbb{B}(R) \Rightarrow Q$;
- on a uo constraint, for some R, $P \Rightarrow R$ and $\mathbb{U}(R) \Rightarrow Q$.

Lo parallelism is defined in definition 5. Embroidery must be stable against lo parallel (internal) interference by rule 4.

An assignment a is *lo parallel* with a constraint b→c if tbere is an so path connecting a, b and c in which a is not constrained to come before b or after c.

Constraint embroidery P is LO stable against lo-parallel assignment A with precondition Q if

$$\mathrm{sp}(P \wedge Q,\, A) \Rightarrow P$$

Embroidery must be stable against external interference by rules 7 and 9.

Constraint embroidery P is EXT stable against $Q \mid x := E$ from another thread if

$$\mathrm{sp}(P \wedge \widehat{Q},\, x := E) \Rightarrow P$$

Constraint embroidery P is UEXT stable against $Q \mid x := E$ from another thread if

$$\mathrm{sp}(\widetilde{P} \wedge Q,\, x := E) \Rightarrow \widetilde{P}$$

Hatting (and double-hatting) is defined in definitions 10, 13, 15, 17, 20 and 23.

$$\begin{array}{rcll}
\widehat{var} & = & \widehat{var} & \text{whether free or bound} \\
\widehat{reg} & = & reg & \\
\widehat{const} & = & const & \\
\widehat{P \;\mathsf{binop}\; Q} & = & \widehat{P}\;\mathsf{binop}\;\widehat{Q} & \\
\widehat{\mathsf{unop}\; P} & = & \mathsf{unop}\;\widehat{P} & \\
\widehat{\mathsf{binder}\; v(P)} & = & \mathsf{binder}\; v(\widehat{P}) & \\
\widehat{\mathbb{B}(P)} & = & \mathbb{B}(P) \wedge \widehat{P} & \\
\widehat{\mathbb{U}(P)} & = & \mathbb{U}(P) \wedge \widehat{P} & \\
\widehat{P \,\mathsf{since}\, Q} & = & \widehat{P}\,\mathsf{since}\,\widehat{Q} & \\
\widehat{\mathbb{S}\mathsf{ofar}(P)} & = & \mathbb{S}\mathsf{ofar}(\widehat{P}) & \\
\widehat{\mathbb{O}\mathsf{uat}(P)} & = & \mathbb{O}\mathsf{uat}(\widehat{P}) & \\
\widehat{v_c(A,B)} & = & v_c(\widehat{A},\widehat{B}) &
\end{array}$$

Twiddling (and double-twiddling) is the same as hatting except for $\mathbb{B}$ (definition 18).

$$\widetilde{\mathbb{B}(P)} \;=\; \widetilde{\mathbb{B}(P)}$$

Stability against quotiented interference is tested by rule 6.

P is stable against $[\overline{ns}].\, Q \mid x := E$ if it is stable against $Q[\overline{ns}\backslash\overline{fs}] \mid x := E[\overline{ns}\backslash\overline{fs}]$ where $\overline{fs}$ are fresh names. Similarly for internal interference $[\overline{ns}].\, Q \mid r := E$ and $[\overline{ns}].\, Q \mid r := x$.

The effect of strongest-post substitutions on modalities is defined in definitions 3, 13, 15, 17, and 20 – all the same except for $\mathbb{O}\mathsf{uat}$. Its effect on coherence assertions is described in definition 23.

$$\begin{array}{rcl}
\mathbb{M}\mathsf{odality}(P)[x\backslash\overleftarrow{x}] & = & \overleftarrow{\mathbb{M}\mathsf{odality}(P)} \wedge P[x\backslash\overleftarrow{x}] \\
\overleftarrow{\mathbb{M}\mathsf{odality}(P)} \wedge P & \Rightarrow & \mathbb{M}\mathsf{odality}(P) \\
\mathbb{M}\mathsf{odality}(P)[r\backslash\overleftarrow{r}] & = & \mathbb{M}\mathsf{odality}(P[r\backslash\overleftarrow{r}]) \\
\mathbb{O}\mathsf{uat}(P)[x\backslash\overleftarrow{x}] & = & \overleftarrow{\mathbb{O}\mathsf{uat}(P)} \\
\overleftarrow{\mathbb{O}\mathsf{uat}(P)} \vee P & \Leftrightarrow & \mathbb{O}\mathsf{uat}(P) \\
\neg\overleftarrow{\mathbb{O}\mathsf{uat}(P)} \wedge \neg P & \Leftrightarrow & \neg\mathbb{O}\mathsf{uat}(P) \\
\mathbb{O}\mathsf{uat}(P)[r\backslash\overleftarrow{r}] & = & \mathbb{O}\mathsf{uat}(P[r\backslash\overleftarrow{r}]) \\
(v_c(A,B))[x\backslash\overleftarrow{x}] & = & v_c(A[x\backslash\overleftarrow{x}], B[x\backslash\overleftarrow{x}]) \\
(v_c(A,B))[r\backslash\overleftarrow{r}] & = & v_c(A[r\backslash\overleftarrow{x}], B[x\backslash\overleftarrow{r}])
\end{array}$$

The rely of a thread is the union of the guarantees of all other threads. In forming the rely we must check BO stability with rule 8 between all interferences not from the same thread and not affecting the same variable, and UO stability between all interferences not from the same thread.

Bo parallelism is defined in definition 6 and BO stability in rule 8.

If variable assignments to distinct variables are so- but not bo-ordered then the so-later is *bo parallel* with the so-earlier and will interfere with its interference precondition.

Interference $P \mid y := F$ is BO stable against bo-parallel $Q \mid x := E$ if

$$\mathrm{sp}(\widehat{P} \wedge \widehat{\widehat{Q}},\ x := E) \Rightarrow \widehat{P}$$

Uo parallelism is defined in definition 19 and UO stability in rule 10.

If variable assignments are so- but not uo-ordered then the so-later is *uo parallel* with the so-earlier and will interfere with the so-earlier's interference precondition. A variable assignment is also uo parallel with itself.

Interference $P \mid x := E$ is UO stable against uo-parallel $Q \mid y := F$ if

$$\mathrm{sp}(\widetilde{P} \wedge \widetilde{\widetilde{Q}},\ y := F) \Rightarrow \widetilde{P}$$

Inclusion in a guarantee is tested by rule 11 with the support of definition 28.

Interference $Q \mid x := E$ is included in the guarantee guar$[g1;\ldots;gn]$ if

$$\mathrm{intf}(Q \mid x := E) \Rightarrow \mathrm{intf}(g1) \vee \ldots \vee \mathrm{intf}(gn) \vee (v1{=}v1' \wedge \ldots \wedge vn{=}vn')$$

where $v1,\ldots,vn{=}\overline{vs}$, the free variables of the guarantee $\bigcup$ free variables of $Q \bigcup \{x\}$.

$$\mathrm{intf}(Q \mid x := E) \triangleq Q \wedge x'{=}E \wedge v1{=}v1' \wedge \ldots \wedge vn{=}vn'$$

where $v1,\ldots,vn{=}\overline{vs}\backslash x$

where $\overline{vs}$ are the free variables of the guarantee $\bigcup$ free variables of $Q \bigcup \{x\}$.

$\mathbb{B}$ and $\mathbb{U}$ modalities are restricted in use by definition 25.

A propagatable assertion P is one for which P is equivalent to $\downarrow P$. An initial-state assertion must be propagatable. In $\mathbb{B}(P)$ and $\mathbb{U}(P)$, P must be propagatable.

$\downarrow P$ (positive occurrences) and $\Downarrow P$ (negative occurrernces) are defined in definition 24.

$$
\begin{array}{rcll}
\downarrow P & = & P & \text{if } P \text{ does not have multiple free variables} \\
\downarrow(\neg P) & = & \neg \Downarrow P & \\
\downarrow(\mathsf{unop}\ P) & = & \mathsf{unop}\ \downarrow P & \\
\downarrow(P\ \mathsf{binop}\ Q) & = & \downarrow P\ \mathsf{binop}\ \downarrow Q & \\
\downarrow(\forall v(P)) & = & \forall v(\downarrow P) & \text{noting that } v \text{ is a bound variable} \\
\downarrow(\exists v(P)) & = & \exists v(\downarrow P) & \text{noting that } v \text{ is a bound variable} \\
\downarrow\mathbb{B}(P) & = & \mathbb{B}(\downarrow P) & \\
\downarrow\mathbb{U}(P) & = & \mathbb{U}(\downarrow P) & \\
\downarrow\mathbb{O}\mathsf{uat}(P\wedge Q) & = & \downarrow\mathbb{O}\mathsf{uat}(P)\wedge\downarrow\mathbb{O}\mathsf{uat}(Q) & \\
\downarrow\mathbb{O}\mathsf{uat}(P\vee Q) & = & \downarrow\mathbb{O}\mathsf{uat}(P)\vee\downarrow\mathbb{O}\mathsf{uat}(Q) & \\
\downarrow\mathbb{O}\mathsf{uat}(P) & = & \mathsf{fresh}\ \mathit{boolv} & \\
\downarrow(P\,\mathsf{since}\,Q) & = & \downarrow P\wedge\downarrow\mathbb{O}\mathsf{uat}(Q) & \\
\downarrow\mathbb{S}\mathsf{ofar}(P\wedge Q) & = & \downarrow\mathbb{S}\mathsf{ofar}(P)\wedge\downarrow\mathbb{S}\mathsf{ofar}(Q) & \\
\downarrow\mathbb{S}\mathsf{ofar}(P) & = & P & \\
\Downarrow\mathbb{B}(P) & = & \mathsf{fresh}\ \mathit{boolv} & \\
\Downarrow\mathbb{U}(P) & = & \mathsf{fresh}\ \mathit{boolv} & \\
\Downarrow\mathbb{O}\mathsf{uat}(P) & = & \mathsf{fresh}\ \mathit{boolv} & \\
\Downarrow(P\,\mathsf{since}\,Q) & = & \mathsf{fresh}\ \mathit{boolv} & \\
\Downarrow\mathbb{S}\mathsf{ofar}(P) & = & \mathsf{fresh}\ \mathit{boolv} &
\end{array}
$$

Coherence is defined in definition 22.

$$
\begin{array}{rl}
\text{membership} & x_c(A,B)\Rightarrow \mathrm{cv}(x) \\
\text{irreflexive} & x_c(A,B)\Rightarrow A{\neq}B \\
\text{transitive} & x_c(A,B)\wedge x_c(B,C)\Rightarrow x_c(A,C) \\
\text{antisymmetric} & x_c(A,B)\Rightarrow \neg x_c(B,A) \\
\text{observed} & \mathbb{O}\mathsf{uat}(\mathbb{O}\mathsf{uat}(x{=}A)\wedge x{=}B)\wedge A{\neq}B\wedge \mathrm{cv}(x)\Rightarrow x_c(A,B)
\end{array}
$$

The program-final assertion is checked according to definition 26, with the assistance of the additional coherence axiom in definition 27.

If the final assertions of the threads are $P_0 \ldots P_{n-1}$, in which register occurrences are translated into $(\mathit{threadnum}{:}\,\mathit{reg})$, and the program-final assertion is ω, then the assertion

$$P_0@0\wedge P_1@1\wedge\ldots\wedge P_{n-1}@(n{-}1)\wedge(\downarrow P_0\wedge\downarrow P_1\wedge\ldots\wedge\downarrow P_{n-1})@n\Rightarrow\omega@n$$

$$\text{pms coherence}\quad x_c(A,B)\Rightarrow \mathrm{cv}(x)\wedge x{\neq}A$$

The properties of modalities are set out in definitions 4, 12, 14, 16 and 21. $\mathbb{B}$, $\mathbb{U}$, since and $\mathbb{S}$ofar share some properties

$$
\begin{array}{rcll}
\mathbb{M}\mathsf{odality}(P) & \Rightarrow & P & \\
\mathbb{M}\mathsf{odality}(P \wedge Q) & = & \mathbb{M}\mathsf{odality}(P) \wedge \mathbb{M}\mathsf{odality}(Q) & \\
\mathbb{M}\mathsf{odality}(\mathbb{M}\mathsf{odality}(P)) & = & \mathbb{M}\mathsf{odality}(P) & \\
\mathbb{M}\mathsf{odality}(P \vee Q) & \Leftarrow & \mathbb{M}\mathsf{odality}(P) \vee \mathbb{M}\mathsf{odality}(Q) & \\
\mathbb{M}\mathsf{odality}(P) & \Leftarrow & P & \text{(if } P \text{ doesn't mention variables)} \\
\mathbb{M}\mathsf{odality}(P \vee Q) & \Rightarrow & \mathbb{M}\mathsf{odality}(P) \vee \mathbb{M}\mathsf{odality}(Q) & \text{(if } P \text{ or } Q \text{ don't mention variables)} \\
(P \Rightarrow Q) & \Rightarrow & \mathbb{M}\mathsf{odality}(P) \Rightarrow \mathbb{M}\mathsf{odality}(Q) & \text{(if } P \Rightarrow Q \text{ is a tautology)}
\end{array}
$$

to which they add some of their own:

$$
\begin{array}{rcl}
\mathbb{U}(P) & \Rightarrow & \mathbb{B}(P) \\
\mathbb{U}(\mathbb{B}(P)) & = & \mathbb{U}(P) \\
\mathbb{B}(\mathbb{U}(P)) & = & \mathbb{U}(P) \\
P \wedge Q & \Rightarrow & P \,\mathsf{since}\, Q \\
P \,\mathsf{since}\, Q & \Rightarrow & P \wedge \mathbb{O}\mathsf{uat}(Q) \\
\mathbb{S}\mathsf{ofar}(P) & \Rightarrow & \mathbb{U}(P) \\
\mathbb{S}\mathsf{ofar}(\mathbb{U}(P)) & = & \mathbb{S}\mathsf{ofar}(P) \\
\mathbb{U}(\mathbb{S}\mathsf{ofar}(P)) & = & \mathbb{S}\mathsf{ofar}(P) \\
\mathbb{S}\mathsf{ofar}(P) & \Rightarrow & \widehat{P} \\
\mathbb{S}\mathsf{ofar}(P) & \Rightarrow & \widehat{\widehat{P}}
\end{array}
$$

($\mathbb{S}\mathsf{ofar}(P)$ also implies $\widetilde{P}$ and $\widetilde{\widetilde{P}}$).

$\mathbb{O}$uat is different:

$$
\begin{array}{rcll}
\mathbb{O}\mathsf{uat}(P) & \Leftarrow & P & \\
\mathbb{O}\mathsf{uat}(P \vee Q) & = & \mathbb{O}\mathsf{uat}(P) \vee \mathbb{O}\mathsf{uat}(Q) & \\
\mathbb{O}\mathsf{uat}(\mathbb{O}\mathsf{uat}(P)) & = & \mathbb{O}\mathsf{uat}(P) & \\
\mathbb{O}\mathsf{uat}(P \wedge Q) & \Rightarrow & \mathbb{O}\mathsf{uat}(P) \wedge \mathbb{O}\mathsf{uat}(Q) & \\
\mathbb{O}\mathsf{uat}(P) & \Rightarrow & P & \text{(if } P \text{ doesn't mention variables)} \\
\mathbb{O}\mathsf{uat}(P \wedge Q) & \Leftarrow & \mathbb{O}\mathsf{uat}(P) \wedge \mathbb{O}\mathsf{uat}(Q) & \text{(if } P \text{ or } Q \text{ don't mention variables)} \\
(P \Rightarrow Q) & \Rightarrow & \mathbb{O}\mathsf{uat}(P) \Rightarrow \mathbb{O}\mathsf{uat}(Q) & \text{(if } P \Rightarrow Q \text{ is a tautology)} \\
\mathbb{O}\mathsf{uat}(P) & \Rightarrow & \neg\mathbb{S}\mathsf{ofar}(\neg P) &
\end{array}
$$

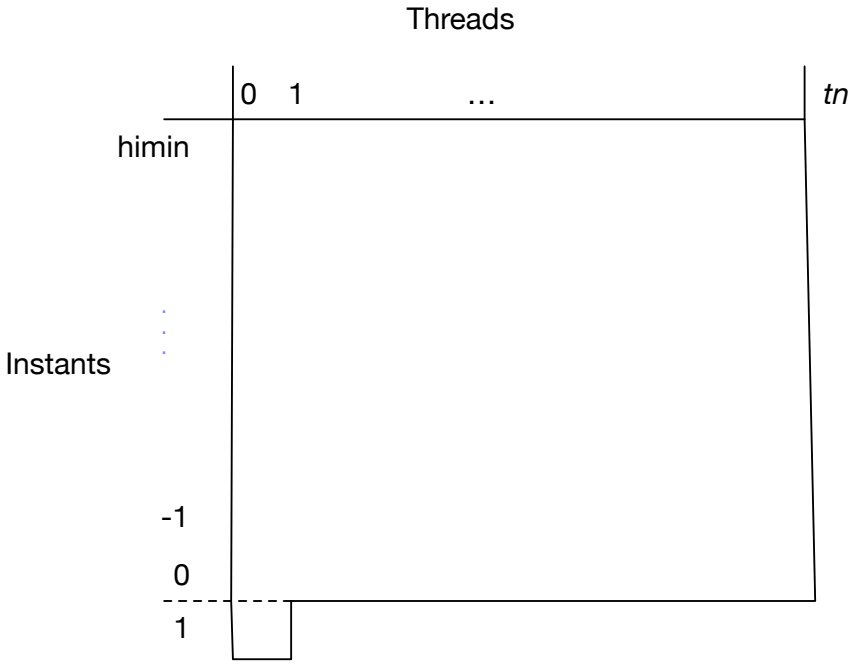

Figure 66: threads × instants

B (S) An SMT embedding

The Arsenic proof checker translates proof obligations into a spatio-temporal logic of threads and events. We use Microsoft Z3 (De Moura and Bjørner, 2008) as the proof-checking oracle.

The semantic domain of the embedding is an array of threads × instants, pictured in fig. 66. Thread numbers $0 \ldots tn-1$ don't correspond to thread numbers in the source program, except in the special case of checking the PMS assertion. The instantaneous effect of an assignment, notionally in thread 0, is to create the special $(0,1)$ element: $(0,0)$ is therefore the 'hooked' state and $(0,1)$ the plain (unhatted, unhooked) state.

- If a query contains hooking and/or hatting and/or twiddling, then *tn* is 2 (threads 0 and 1); with double hatting and/or double twiddling, it is 3; in either case 'now' is state 1 and the initial instant is *himin*.
- Otherwise, if a query contains thread-numbered formulae $P@n$ – which can happen in the PMS assertion – then *tn* is the greatest thread number plus one, and 'now' is 0.
- Otherwise, if a query contains $\mathbb{S}$ofar or $\mathbb{U}$ then *tn* is 2 and 'now' is 0.
- Otherwise, *tn* is 1 and 'now' is 0.

In fig. 67 $[\![P]\!]_I^T$ is the value of P in thread T at instant I. $\mathbb{F}$andw – 'far and wide' – is a modality invented to aid the description of the embedding.

The basis of the embedding is the treatment of variables. $[\![x]\!]_I^T$ is $\mathrm{val}_{\tau(x)}(x,T,I)$, where $\mathrm{val}_{\tau(x)}$ is a function whose result type is the type of x in the source program, and x in the embedded formula is an uninterpreted name. A hatted variable $\widehat{x}$ is embedded as $\mathrm{val}_{\tau(x)}(x,1,hatI)$, where *hatI* is constrained to be negative; similarly $\widehat{\widehat{x}}$ is $\mathrm{val}_{\tau(x)}(x,2,dhatI)$. Since hatting and twiddling never occur together in the same proof obligation, we can use the same embedding for twiddling as for hatting. Hooked and hatted modalities are shifted from $(1,0)$ to $(0,0)$, $(1,hatI)$ and $(1,dhatI)$ just like variables. Register hooking is renaming.

Translation of 'since' is the basis of our embedding of modalities. $\mathbb{B}$ and $\mathbb{U}$ are versions of 'since' which use a special 'boundary event' variable *bev*. To provide $\mathbb{S}\mathsf{ofar}(P) \Rightarrow \mathbb{B}(P)$, $\mathbb{S}\mathsf{ofar}(P) \Rightarrow \mathbb{U}(P)$ and $\mathbb{B}(\text{true}) = \mathbb{U}(\text{true}) = \text{true}$ the embedding ensures that *bev* holds at *himin*, a state constrained to be before any other state.

Coherence assertions are translated very like variables; cv assertions are unchanged.

Because state $(0,1)$, and state $(0,0)$, from which it derives, use different variable embeddings, sp requires careful treatment. The definition of sp and variable assignment that the proof-checker uses makes it clear that only the assigned variable changes value.

Definition 29: Strongest post sp in the presence of temporal assertions

$$\mathrm{sp}(P,\ x := E) \triangleq P[x \backslash \overleftarrow{x}] \wedge x{=}E \wedge \overleftarrow{y}{=}y \wedge \overleftarrow{z}{=}z \wedge \ \ldots$$

where $y, z, \ldots$ are the variables, other than x, free in P.

To reduce the stress on Z3 caused by alternation of $\exists$ and $\forall$ quantifiers we employ some equivalences, shown in fig. 68, many of which derive from definitions 12, 14, 16 and 21.

$$
\begin{array}{rcll}
[\![x]\!]^T_I & = & \mathrm{val}_{\tau(x)}(x,T,I) & \\
[\![r]\!]^T_I & = & r & \\
[\![\mathrm{constant}]\!]^T_I & = & \mathrm{constant} & \\
[\![\overleftarrow{r}]\!]^0_0 & = & r' & \\
[\![\overleftarrow{P}]\!]^0_1 & = & [\![P]\!]^0_0 & \text{variables and modalities} \\
[\![\widehat{P}]\!]^0_1 & = & [\![P]\!]^1_{hatI} & \text{variables and modalities} \\
[\![\widehat{\widehat{P}}]\!]^0_1 & = & [\![P]\!]^2_{dhatI} & \text{variables and modalities} \\
[\![\tilde{P}]\!]^0_1 & = & [\![P]\!]^1_{hatI} & \text{variables and modalities} \\
[\![\widetilde{\widetilde{P}}]\!]^0_1 & = & [\![P]\!]^2_{dhatI} & \text{variables and modalities} \\
[\![P\,\mathsf{since}\,Q]\!]^T_I & = & \exists j \left(himin \le j \le I \wedge [\![Q]\!]^T_j \wedge \forall j' \left(j \le j' \le I \Rightarrow [\![P]\!]^T_{j'} \right) \right) & \\
[\![\mathbb{B}(P)]\!]^T_I & = & [\![P\,\mathsf{since}\,bev]\!]^T_I & \\
[\![\mathbb{U}(P)]\!]^T_I & = & [\![\mathbb{F}\mathsf{andw}(P)\,\mathsf{since}\,bev]\!]^T_I & \\
[\![\mathbb{S}\mathsf{ofar}(P)]\!]^T_I & = & \forall j (himin \le j \le I \Rightarrow [\![\mathbb{F}\mathsf{andw}(P)]\!]^T_j) & \\
[\![\mathbb{O}\mathsf{uat}(P)]\!]^T_I & = & \exists j (himin \le j \le I \wedge [\![P]\!]^T_j) & \\
[\![\mathbb{F}\mathsf{andw}(P)]\!]^0_1 & = & [\![P]\!]^0_1 & \\
[\![\mathbb{F}\mathsf{andw}(P)]\!]^T_I & = & \forall t (0 \le t < tn \Rightarrow [\![P]\!]^t_I) & \\
[\![\mathsf{binder}\,x(P)]\!]^T_I & = & \mathsf{binder}\,x([\![P]\!]^T_I) & \\
[\![P\ \mathsf{binop}\ Q]\!]^T_I & = & [\![P]\!]^T_I\ \mathsf{binop}\ [\![Q]\!]^T_I & \\
[\![\mathsf{unop}\ P]\!]^T_I & = & \mathsf{unop}\ [\![P]\!]^T_I & \\
[\![P@n]\!]^T_0 & = & [\![P]\!]^n_0 & \\
[\![x_c(A,B)]\!]^T_I & = & \mathrm{co}_{\tau(x)}(x, [\![A]\!]^T_I, [\![B]\!]^T_I) & \\
[\![\mathrm{cv}(x)]\!]^T_I & = & \mathrm{cv}(x) &
\end{array}
$$

Figure 67: Rules of embedding

$$
\begin{array}{rcl}
[\![(P \,\mathsf{since}\, Q) \,\mathsf{since}\, R]\!]_I^T & = & [\![P \,\mathsf{since}\, ((P \,\mathsf{since}\, Q) \wedge R)]\!]_I^T \\
[\![\mathbb{S}\mathsf{ofar}(P) \,\mathsf{since}\, Q]\!]_I^T & = & [\![\mathbb{S}\mathsf{ofar}(P) \wedge \mathbb{O}\mathsf{uat}(Q)]\!]_I^T \\
[\![\mathbb{U}(\mathbb{U}(P))]\!]_I^T & = & [\![\mathbb{U}(P)]\!]_I^T \\
[\![\mathbb{U}(\mathbb{B}(P))]\!]_I^T & = & [\![\mathbb{U}(P)]\!]_I^T \\
[\![\mathbb{U}(\mathbb{S}\mathsf{ofar}(P))]\!]_I^T & = & [\![\mathbb{S}\mathsf{ofar}(P)]\!]_I^T \\
[\![\mathbb{U}(P \,\mathsf{since}\, Q)]\!]_I^T & = & [\![\mathbb{F}\mathsf{andw}(P) \,\mathsf{since}\, (\mathbb{F}\mathsf{andw}(P) \wedge Q)]\!]_I^T \\
[\![\mathbb{S}\mathsf{ofar}(\mathbb{U}(P))]\!]_I^T & = & [\![\mathbb{S}\mathsf{ofar}(P)]\!]_I^T \\
[\![\mathbb{S}\mathsf{ofar}(\mathbb{B}(P))]\!]_I^T & = & [\![\mathbb{S}\mathsf{ofar}(P)]\!]_I^T \\
[\![\mathbb{S}\mathsf{ofar}(\mathbb{S}\mathsf{ofar}(P))]\!]_I^T & = & [\![\mathbb{S}\mathsf{ofar}(P)]\!]_I^T
\end{array}
$$

Figure 68: Embedding equivalences